



Learning Generative Factors of EEG Data with Variational Auto-Encoders

Maksim Zhdanov^{1,2(✉)}, Saskia Steinmann³, and Nico Hoffmann¹

¹ Helmholtz-Zentrum Dresden-Rossendorf, Dresden, Germany
maxxxzdn@gmail.com

² TU Dresden, Dresden, Germany

³ Psychiatry Neuroimaging Branch, Department of Psychiatry and Psychotherapy,
University Medical Center Hamburg-Eppendorf, Hamburg, Germany

Abstract. Electroencephalography produces high-dimensional, stochastic data from which it might be challenging to extract high-level knowledge about the phenomena of interest. We address this challenge by applying the framework of variational auto-encoders to 1) classify multiple pathologies and 2) recover the neurological mechanisms of those pathologies in a data-driven manner. Our framework learns generative factors of data related to pathologies. We provide an algorithm to decode those factors further and discover how different pathologies affect observed data. We illustrate the applicability of the proposed approach to identifying schizophrenia, either followed or not by auditory verbal hallucinations. We further demonstrate the ability of the framework to learn disease-related mechanisms consistent with current domain knowledge. We also compare the proposed framework with several benchmark approaches and indicate its classification performance and interpretability advantages.

Keywords: EEG · VAEs · Functional connectivity

1 Introduction

Analysis of neurological processes in the human brain is a challenging process addressed by neuroimaging. Here, one typically obtains high-dimensional stochastic data, which encourages the usage of machine learning algorithms. In recent years, deep learning discriminative models have been actively applied to neuroimaging issues (see [1] for a review). They yielded state-of-the-art results in classification problems on a variety of benchmark datasets [2–4]. One downside of deep learning-based classifiers is that they operate as black boxes [5] meaning that interpreting their predictions is often severely complicated.

S. Steinmann and N. Hoffmann—Equal contribution.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/978-3-031-18576-2_5.

[1] suggested that hybrid generative-discriminative models might help resolve the issue. Such models can learn low-dimensional representations of data where each dimension corresponds to an independent *generative factor* (i.e. a disentangled representation, see [6] for a review). The discriminative part of the model then forces those factors to capture label information from data [7]. *Interpretability* is thus achieved via decoding the meaning of generative factors related to particular labels [8]. It is especially relevant in neuroimaging as one can observe how underlying pathologies govern the process of data generation.

The paper follows the intuition regarding hybrid generative-discriminative models for neuroimaging data, with particular application to EEG data. Our main contributions are as follows:

1. We demonstrate how one can apply characteristic capturing variational auto-encoders (CCVAEs) [7] to the interpretable classification of EEG data;
2. We compare the model to two generative models previously used for EEG data: conditional VAEs and VAEs with downstream classification;
3. We propose an algorithm for decoding generative factors learned by CCVAEs;
4. We demonstrated that learned generative mechanisms associated with pathologies are consistent with evidence from neurobiological studies.

2 Background

In this section, we introduce the relevant materials on variational auto-encoders, disentangled factorization and the role of supervision.

Variational Auto-Encoders. Variational auto-encoders (VAEs) [9] learn a model distribution $p_\theta(\mathbf{x}, \mathbf{z})$ that describes the ground-truth data generation process $p(\mathbf{x}, \mathbf{z})$ as first sampling random variables \mathbf{z} from a prior distribution $p(\mathbf{z})$. Then, an observation \mathbf{x} is inferred based on generative factors $p_\theta(\mathbf{x}|\mathbf{z})$ yielding

$$p_\theta(\mathbf{x}, \mathbf{z}) = p_\theta(\mathbf{x}|\mathbf{z})p(\mathbf{z}) \quad (1)$$

Here, the conditional distribution is parameterized with neural networks whose learned parameters are denoted with θ . Defining latent variables as jointly independent yields *disentangled factorization* [10] that separates the generative process into human-interpretable [8] generative mechanisms.

Supervised Learning. A label variable $\mathbf{y} \sim p(\mathbf{y})$ can be interpreted as the context that partially governs the generation of an observed variable \mathbf{x} . In VAEs, it is reflected by generative factors $p(\mathbf{x}|\mathbf{z}, \mathbf{y})$ of a model. It leads to the joint distribution factorized as follows:

$$p_{\theta_1, \theta_2}(\mathbf{x}, \mathbf{z}, \mathbf{y}) = p_{\theta_1}(\mathbf{x}|\mathbf{z}, \mathbf{y})p_{\theta_2}(\mathbf{z}|\mathbf{y})p(\mathbf{y}) \quad (2)$$

where θ_1, θ_2 are parameters of corresponding model distributions. The equality holds due to the chain rule. Incorporating label information into the model allows learning generative factors corresponding to those labels via $p_{\theta_2}(\mathbf{z}|\mathbf{y})$.




Fig. 1. Scheme of the proposed approach (A). We receive EEG data as input and learn a stochastic mapping to the latent space with CCVAEs (B) [7]. We further manipulate learned generative factors of data to gain insights regarding neurological mechanisms related to the attribute of interest, e.g. a symptom.

3 Methods

The proposed framework consists of 2 steps (see Fig. 1). First, EEG data is mapped stochastically to the latent space via CCVAEs. The latent space is constructed such that each label is related to a single independent generative factor. Second, we perform an intervention analysis to decode the meaning of label-related generative factors. This way, we get an intuition regarding mechanisms through which labels govern data generation. In our case, we are interested how different pathologies manifest themselves in functional connectivity matrices.

Characteristic Capturing VAEs. We aim at learning a model of a joint distribution over observed EEG data \mathbf{x} , labels (e.g. pathology indicators) \mathbf{y} and latent variables \mathbf{z} partially conditioned to \mathbf{y} . Let us assume that \mathbf{x} and \mathbf{y} are conditionally independent given \mathbf{z} . Then, the generative model (see Eq. 2) can be rewritten as follows:

$$p_{\theta_1, \theta_2}(\mathbf{x}, \mathbf{y}, \mathbf{z}) = p_{\theta_1}(\mathbf{x}|\mathbf{z}) p_{\theta_2}(\mathbf{z}|\mathbf{y}) p(\mathbf{y})$$

We further partition the latent space \mathbf{z} such that one partition \mathbf{z}_c encapsulates label associated characteristics, and the second partition $\mathbf{z}_{\setminus c}$ accounts for shared features of data (as in the vanilla VAEs):

$$p_{\theta_2}(\mathbf{z}|\mathbf{y}) = p_{\theta_2}(\mathbf{z}_c|\mathbf{y}) \cdot p(\mathbf{z}_{\setminus c})$$

The characteristic partition \mathbf{z}_c is further partitioned so that each label can access only a single latent variable. It guarantees the disentanglement of label information in latent representations. The intractable distribution $p(\mathbf{z}|\mathbf{x}, \mathbf{y})$ is conditioned to both observation and label variables. It is approximated with the following inference model:

$$q_{\phi_1, \phi_2}(\mathbf{z}|\mathbf{x}, \mathbf{y}) = \frac{q_{\phi_1}(\mathbf{y}|\mathbf{z}_c) q_{\phi_2}(\mathbf{z}|\mathbf{x})}{q_{\phi_1, \phi_2}(\mathbf{y}|\mathbf{x})}$$

where ϕ_1, ϕ_2 are parameters of model distributions. The conditional distribution

$$q_{\phi_1, \phi_2}(\mathbf{y}|\mathbf{x}) = \int q_{\phi_1}(\mathbf{y}|\mathbf{z}_c) q_{\phi_2}(\mathbf{z}|\mathbf{x}) d\mathbf{z}$$

reflects that observation variables \mathbf{x} and label variables \mathbf{y} are connected via the characteristic partition \mathbf{z}_c . Label-related information is captured in an observation \mathbf{x} by the inference model $q_{\phi_2}(\mathbf{z}|\mathbf{x})$. At the same time, classifier $q_{\phi_1}(\mathbf{y}|\mathbf{z}_c)$ forces the label-related latent variables \mathbf{z}_c to capture characteristics of those labels.

As for the vanilla VAEs, the model is optimized by maximizing the evidence lower bound [9]. In the case of CCVAEs, it is equivalent to maximizing the following objective (see Appendix B.1 of [7] for derivation):

$$\mathcal{L}(\mathbf{x}, \mathbf{y}) = \mathbb{E}_{q_{\phi_2}(\mathbf{z}|\mathbf{x})} \left[\frac{q_{\phi_1}(\mathbf{y}|\mathbf{z}_c)}{q_{\phi_1, \phi_2}(\mathbf{y}|\mathbf{x})} \log \frac{p_{\theta_1}(\mathbf{x}|\mathbf{z}) p_{\theta_2}(\mathbf{z}|\mathbf{y})}{q_{\phi_1}(\mathbf{y}|\mathbf{z}_c) q_{\phi_2}(\mathbf{z}|\mathbf{x})} \right] + \log q_{\phi_1, \phi_2}(\mathbf{y}|\mathbf{x}) \quad (3)$$

The classification term $\log q_{\phi_1, \phi_2}(\mathbf{y}|\mathbf{x})$ is essentially a learnable mapping from input data \mathbf{x} to labels \mathbf{y} that goes through the characteristic partition of the latent space \mathbf{z}_c . It applies pressure onto the partition to learn label-related characteristics from data and simultaneously performs data classification.

Intervention Analysis. The learned generative model forms the bridge between observations \mathbf{x} and their labels \mathbf{y} via latent variables \mathbf{z}_c . It allows one to analyze generative factors $p_{\theta_1, \theta_2}(\mathbf{x}|\mathbf{z}, \mathbf{y})$ of data related to those labels. One can explore the relation via intervention analysis. The algorithm for a single *binary* label of interest \mathbf{y}^i is as follows. First, one fixes every dimension of the latent space \mathbf{z} except the one \mathbf{z}_c^i that corresponds to the label \mathbf{y}^i . Next, the value of \mathbf{z}_c^i is sampled from $p_{\theta_2}(\mathbf{z}_c^i|\mathbf{y}^i)$ for each value of $\mathbf{y}^i \in \{0, 1\}$. As a result, one receives two latent representations $\mathbf{z}_0, \mathbf{z}_1$ that vary only in a single dimension \mathbf{z}_c^i . Those representations are then reconstructed to the observation space $\mathbf{x}_0 \sim p_{\theta_1}(\mathbf{x}|\mathbf{z} = \mathbf{z}_0), \mathbf{x}_1 \sim p_{\theta_1}(\mathbf{x}|\mathbf{z} = \mathbf{z}_1)$. The procedure is repeated for N times. As a result, one gets multiple pairs of reconstructions $(\mathbf{x}_0, \mathbf{x}_1)$ that are different only to the varied generative factor \mathbf{z}_c^i . One further calculates the average difference $\frac{1}{N} \sum_{k=1}^N (\mathbf{x}_1^k - \mathbf{x}_0^k)$ for each pair, and thus observes how the label \mathbf{y}^i manifests itself in data.

4 Related Works

The fusion of generative and discriminative models with application to neuroimaging data is an active area of research. [11] demonstrate that using learned representations leads to more robust classification performance compared to feed-forward neural networks. [12] introduce VAEs into feature extraction from multichannel EEG data yielding better accuracy than traditional unsupervised

approaches. [13] use stacked VAEs for semi-supervised learning on EEG data. However, the label information is usually encapsulated by multiple latent variables simultaneously. In this case, label characteristics are smeared across the latent space, thus complicating the analysis of label-related generative factors. It, in turn, limits both the interpretability and explainability of these models. One has to decode and interpret each latent variable and then infer the relation with label variables which is not a trivial task.

Two flavours of VAEs that are commonly applied to EEG data are conditional VAEs [13]¹ and VAEs with downstream classification [12]. In both approaches, the latent space is not partitioned with respect to label variables. Hence, compared to CCVAEs, their general disadvantage is reduced interpretability of classification as it is difficult to build a bridge between labels and generative factors.

Conditional VAEs. Conditional VAEs have a graphical model similar to the one of CCVAEs. The only difference is that the latent space is not partitioned to labels, i.e. $\mathbf{z} = \mathbf{z}_c$. Learnable parameters are optimized via maximizing the objective Eq. (3). The framework allows conditional sampling, so one can use intervention analysis to decode the meaning of learned generative factors. Nevertheless, the interpretation is complicated as a single label variable is connected to each dimension of the latent space.

VAEs + Downstream Classification. The model approximates the joint distribution of observed data and latent variables that is factorized as Eq. (1). The relation between latent variables and labels is built via classifying a latent representation. The model is trained via optimizing the following objective [11]:

$$\mathcal{L}(\mathbf{x}, \mathbf{y}) = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} \left[\log p_\theta(x|z) - D_{KL}(q_\phi(z|x)||p(z)) - BCE(f_\xi(z), y) \right] \quad (4)$$

where $f_\xi : Z \rightarrow Y$ is a learnable classifier with parameters ξ , BCE is binary cross-entropy function. Here, the information about label variables is incorporated into the latent space via pressure applied by a downstream classification task. The model can be seen as a feed-forward deep neural network with additional regularization imposed by the decoder part of VAEs.

5 Experimental Details

Experimental Study. The study comprised 29 patients suffering from schizophrenia and 52 healthy controls. 14 subjects out of those 29 indicated the emergence of auditory verbal hallucinations (AVH), i.e. hearing voices with no external stimuli presented. Every participant was right-handed. Six different

¹ Technically, [13] use stacked VAEs that have two connected latent spaces. One of the spaces is connected to label variables. However, the framework can be seen as an instance of conditional VAEs with a non-trivial structure of the latent space.

syllables were spoken to each participant (/ba/, /da/, /ka/, /ga/, /pa/, /ta/) for 500 ms simultaneously to each ear after 200 ms silence period. Meanwhile, the EEG recording was conducted with 64 electrodes where 4 EOG channels were used to monitor eye movements. For each subject, we repeated the procedure multiple times (number of trials for AVH: 68.23 ± 19.43 ; SZ: 68.76 ± 14.79 and HC: 71.19 ± 12.93). At the preprocessing step, the data was filtered from 20 to 120 Hz according to a protocol described in [14]. Therefore, only gamma-band frequencies are preserved. Afterwards, all channels were re-referenced to the common average. At last, muscle and visual artefacts were identified and removed. For our experiments, we utilized two parts of a recording: the resting one (first 200 ms with no syllable given) and the listening one (initial 200 ms when syllables were presented). The study of [14] contains detailed data acquisition and preprocessing information.

Experimental Data. For each EEG recording $[\zeta_1, \zeta_2, \dots, \zeta_{61}]$, we assessed functional connectivity by calculating a correlation matrix:

$$\mathbf{x}_{ij} = \frac{\text{cov}(\zeta_i, \zeta_j)}{\sqrt{\text{var}(\zeta_i) \cdot \text{var}(\zeta_j)}}$$

As a result, functional connectivity matrices play the role of observed data \mathbf{x} . We introduce 3 binary labels such that $\mathbf{y} = [\textit{listening}, \textit{schizophrenia}, \textit{hallucinations}]$. To create a dataset for training models, we use the intra-patient paradigm, i.e. data from the same subject can appear simultaneously in training and test datasets. Thus, data of all the subjects are randomly sampled to form those datasets yielding 9000 training samples and 2000 test samples.

Implementation Details. One can find details regarding the parametrization of distributions in the supplementary material (Section S.1). We release the implementation at [GitHub](#). For each framework, the parameters θ_i, ϕ_j (and ξ for VAEs) are trained via optimizing the corresponding objective. We use Adam optimizer with a learning rate of 10^{-3} . The training was performed in mini-batches of size 32 for 100 epochs. All models are trained on an NVIDIA Tesla V100 GPU from the Hemera HPC system of HZDR.

6 Results and Discussion

We found that high-dimensional latent spaces ($\text{dim} > 32$) hinder the reproducibility of generative factors learned by CCVAEs. For that reason, we keep the latent space of all models low-dimensional: $\mathbf{z} \in \mathbb{R}^5$ ($\mathbf{z}_c \in \mathbb{R}^3, \mathbf{z}_{\setminus c} \in \mathbb{R}^2$ for CCVAEs).

Results. As shown in Table 1, CCVAEs outperform baseline models in both classification performance and disentanglement (see supplementary material S.2 for details). The framework consistently classifies observed data based

on its low-dimensional representation, yielding a low standard deviation of accuracy. Besides, it demonstrates a high level of disentanglement, meaning that each label variable is captured only by a single latent dimension. For CCVAEs, generative factors are disentangled in the latent space by design, leading to the highest score. It is not as high as expected due to the correlation between pathology labels.




Fig. 2. Confusion matrices for different methods. The size of the circle indicates the value of the corresponding element. Rows correspond to label variables (L - *listening*, S - *schizophrenia*, H - *hallucinations*) while columns represent latent generative factors.

Table 1. Comparison of CCVAEs to baseline models in terms of accuracy and disentanglement scores on the test dataset. For each framework, 10 experiments were conducted.

Framework	Accuracy	SAP score [16]	MIG score [15]
CCVAEs	0.84 ± 0.01	0.34 ± 0.02	0.04 ± 0.01
Conditional VAEs	0.69 ± 0.10	0.07 ± 0.07	0.01 ± 0.01
VAEs + classification	0.74 ± 0.03	0.06 ± 0.04	0.01 ± 0.01




Fig. 3. Latent space generated by sampling from the inference model $q_\phi(\mathbf{z}|\mathbf{x})$ of different methods. For CCVAEs, the axes corresponding to pathology variables are shown. For baseline methods, 2 randomly selected dimensions are visualized.

Disentangled Latent Space. To demonstrate how hard-wired disentanglement affects the latent space learned by CCVAEs, we construct confusion matrices in the following way. We compute latent representation for each data point in the test dataset and intervene (i.e. randomly change its value) upon a single dimension. We further observe how log-probabilities assigned by a pre-trained classifier change due to the intervention for each label. We calculate the difference for each label-latent pair yielding a confusion matrix (see Fig. 2). The optimal result would be one non-zero element per row (i.e. label), which means that each label corresponds to only a single generative factor. This is the case of CCVAEs, where one can observe one-to-one dependence between label variables and corresponding generative factors. This leads to latent representations being robust to variations in data generative factors, as those are independent by design. At the same time, the characteristics of labels are entangled within latent spaces of baseline frameworks. Hence, it is difficult to disentangle the influence of a label from other generative factors, which severely hinders interpretability.

Posterior Distribution. We further compare latent spaces learned by each framework. To visualize the latent space, we sample $\mathbf{z} \sim q_\phi(\mathbf{z}|\mathbf{x})$ for multiple \mathbf{x} from the test dataset for each model. The result is shown in Fig. 3. In the case of CCVAEs, the distribution has three modes corresponding to subject cohorts in data (healthy, schizophrenia, schizophrenia followed by AVH). The separation is caused by the influence of the conditional prior and the classifier, aiming to separate representations encoding different label combinations. In the case of baseline frameworks, there is no strict regularisation that preserves label information within a partition of the latent space. As a result, features of data encoded by their latents are shared between cohorts of subjects (thus one or two modes). We discovered that both baseline methods often fail to jointly learn a low-dimensional representation and classify labels when the pressure on the KL divergence term in the loss objective is high. The problem is partially solved by introducing a scaling factor β for the term [8]. However, reducing the pressure might lead to untrustable reconstruction if prior $p(\mathbf{z})$ is not sufficiently close to the inference model $q(\mathbf{z}|\mathbf{x})$. This is not the case for CCVAEs that do not require any manual fine-tuning and operate stably with low-dimensional latent spaces.

Analyzing Pathological Mechanisms.

We further investigate what connections are affected when intervening upon a single label dimension via intervention analysis (Fig. 4, see supplementary material S.3 for computation details). The model associates the emergence of AVH with alterations in frontotemporal brain areas (the highest positive difference), which have been repeatedly observed in prior studies [17,18]. The salient connections are mainly located in the right hemisphere, which is supported by the fMRI study of [19]. The model also points toward reduced connectivity between hemispheres. It is coherent with the current hypothesis (see [14] for review) that connects the emergence of auditory verbal hallucinations with the interhemispheric miscommunication during auditory processing. Overall, the model can at least partially reconstruct the neurological mechanism of the symptom for functional connectivity. To explain the emergence of schizophrenia, the model focuses mainly on the left hemisphere. It is not surprising since the auditory function is left-lateralized for right-handed people [20,21]. It would be an interesting direction for further studies to apply CCVAEs to learn the mechanisms of particular symptoms of the composite disorder (e.g. hallucinations, delusions, etc.).




Fig. 4. Average difference in reconstructions of functional connectivity matrices when intervening on a single label: *schizophrenia* (Left) and *hallucinations* (Right). Connections that are stronger when a disorder is presented are shown in red; otherwise, blue. For clarity, we visualize only 40 connections with the highest absolute value.

7 Conclusion

We demonstrated how to apply the framework of characteristic capturing variational auto-encoders to EEG data analysis. The method encapsulates and disentangles the characteristics associated with different pathologies in the latent space. As generative factors are independent by design, one can decode their meaning and discover how those pathologies alter observed data. It leads to improved interpretability coupled with the high classification performance of neural networks. The framework is not limited to functional connectivity analysis or EEG data and can be easily adapted to different neuroimaging modalities.

References

1. Davatzikos, C.: Machine learning in neuroimaging: progress and challenges. *Neuroimage* **197**, 652–656 (2019)
2. Zhang, L., Wang, M., Liu, M., Zhang, D.: A survey on deep learning for neuroimaging-based brain disorder analysis. *Front. Neurosci.* **14**, 779 (2020)
3. Wahlang, I., et al.: Brain magnetic resonance imaging classification using deep learning architectures with gender and age. *Sensors* **22**, 1766 (2022)
4. Li, Y., et al.: A novel bi-hemispheric discrepancy model for EEG emotion recognition. *IEEE Trans. Cogn. Dev. Syst.* **13**, 354–367 (2021)
5. Quinn, T.P., Jacobs, S., Senadeera, M., Le, V., Coghlan, S.: The three ghosts of medical AI: can the black-box present deliver? *Artif. Intell. Med.* **124**, 102158 (2022)
6. Liu, X., Sanchez, P., Thermos, S., O’Neil, A.Q., Tsafaris, S.A.: Learning disentangled representations in the imaging domain. *Med. Image Anal.* (2022)
7. Joy, T., Schmon, S.M., Torr, P.H., Siddharth, N., Rainforth, T.: Capturing label characteristics in VAEs. In: *ICLR* (2021)
8. Higgins, I., et al.: Beta-VAE: learning basic visual concepts with a constrained variational framework. In: *ICLR* (2017)
9. Kingma, D.P., Welling, M.: Auto-Encoding Variational Bayes. *CoRR*, abs/1312.6114 (2014)
10. Locatello, F., Bauer, S., Lucic, M., Gelly, S., Schölkopf, B., Bachem, O.: Challenging Common Assumptions in the Unsupervised Learning of Disentangled Representations. *arXiv*, abs/1811.12359 (2019)
11. Krishna, G., Tran, C., Carnahan, M., Tewfik, A.H.: Constrained Variational Autoencoder for improving EEG based Speech Recognition Systems. *arXiv*, abs/2006.02902 (2020)
12. Li, X., et al.: Latent factor decoding of multi-channel EEG for emotion recognition through autoencoder-like neural networks. *Front. Neurosci.* **14**, 87 (2020)
13. Chen, J., Yu, Z., Gu, Z.: Semi-supervised Deep Learning in Motor Imagery-Based Brain-Computer Interfaces with Stacked Variational Autoencoder (2020)
14. Steinmann, S., Leicht, G., Andreou, C., Polomac, N., Mulert, C.: Auditory verbal hallucinations related to altered long-range synchrony of gamma-band oscillations. *Sci. Rep.* **7**(1), 1–10 (2017)
15. Chen, T.Q., Li, X., Grosse, R.B., Duvenaud, D.K.: Isolating sources of disentanglement in variational autoencoders. In: *NeurIPS* (2018)
16. Kumar, A., Sattigeri, P., Balakrishnan, A.: Variational Inference of Disentangled Latent Concepts from Unlabeled Observations. *arXiv*, abs/1711.00848 (2018)

17. Jardri, R., Pouchet, A., Pins, D., Thomas, P.: Cortical activations during auditory verbal hallucinations in schizophrenia: a coordinate-based meta-analysis. *Am. J. Psychiatry* **168**(1), 73–81 (2011)
18. Lavigne, K.M., et al.: Left-dominant temporal-frontal hypercoupling in schizophrenia patients with hallucinations during speech perception. *Schizophr. Bull.* **41**(1), 259–67 (2015)
19. Hwang, M., et al.: Auditory hallucinations across the psychosis spectrum: evidence of dysconnectivity involving cerebellar and temporal lobe regions. *NeuroImage Clin.* **32**, 102893 (2021)
20. Papathanassiou, D., Etard, O., Mellet, E., Zago, L., Mazoyer, B., Tzourio-Mazoyer, N.: A common language network for comprehension and production: a contribution to the definition of language epicenters with PET. *Neuroimage* **11**, 347–357 (2000)
21. Flinker, A., et al.: Redefining the role of Broca’s area in speech. *Proc. Natl. Acad. Sci.* **112**, 2871–2875 (2015)