

000 AUTOREGRESSIVE VISUAL DECODING FROM EEG 001 SIGNALS 002

003 **Anonymous authors**
004
005 Paper under double-blind review
006

007 ABSTRACT 008

009
010
011 Electroencephalogram (EEG) signals have become a popular medium for decod-
012 ing visual information due to their cost-effectiveness and high temporal resolution.
013 However, current approaches face significant challenges in bridging the modal-
014 ity gap between EEG and image data. These methods typically rely on complex
015 adaptation processes involving multiple stages, making it hard to maintain con-
016 sistency and manage compounding errors. Furthermore, the computational over-
017 head imposed by large-scale diffusion models limit their practicality in real-world
018 brain-computer interface (BCI) applications. In this work, we present AVDE, a
019 lightweight and efficient framework for visual decoding from EEG signals. First,
020 we leverage LaBraM, a pre-trained EEG model, and fine-tune it via contrastive
021 learning to align EEG and image representations. Second, we adopt an autoregres-
022 sive generative framework based on a "next-scale prediction" strategy: images are
023 encoded into multi-scale token maps using a pre-trained VQ-VAE, and a trans-
024 former is trained to autoregressively predict finer-scale tokens starting from EEG
025 embeddings as the coarsest representation. This design enables coherent genera-
026 tion while preserving a direct connection between the input EEG signals and the
027 reconstructed images. Experiments on two datasets show that AVDE outperforms
028 previous state-of-the-art methods in both image retrieval and reconstruction tasks,
029 while using only 10% of the parameters. In addition, visualization of intermediate
030 outputs shows that the generative process of AVDE reflects the hierarchical nature
031 of human visual perception. These results highlight the potential of autoregressive
032 models as efficient and interpretable tools for practical BCI applications. The code
033 is available at <https://anonymous.4open.science/r/avde-783D>.
034

035 1 INTRODUCTION 036

037 How can we access and interpret the rich visual information encoded in human brain activity? This
038 question has captivated neuroscientists for decades, driving fundamental research at the intersection
039 of cognitive science and artificial intelligence. Decoding human vision from non-invasive neural
040 signals not only advances our understanding of neural representation mechanisms but also promises
041 transformative applications in brain-computer interfaces. Early pioneering work (Kay et al., 2008;
042 Miyawaki et al., 2008; Naselaris et al., 2009) established that simple visual patterns could be de-
043 coded from functional magnetic resonance imaging (fMRI), while recent advances in generative
044 AI have enabled reconstruction of remarkably detailed visual content from brain signals (Takagi &
045 Nishimoto, 2023a; Scotti et al., 2023; Fang et al., 2023).

046 Despite these successes, fMRI-based approaches face fundamental limitations for practical applica-
047 tions: they operate at temporal resolutions orders of magnitude slower than actual neural process-
048 ing, require costly infrastructure, and confine subjects to restrictive scanner environments (Menon &
049 Kim, 1999; Logothetis, 2008). These constraints have motivated a shift toward electroencephalog-
050 raphy (EEG) for visual decoding (Cichy & Pantazis, 2017). EEG offers millisecond-level temporal
051 precision, while providing significantly greater portability and accessibility at a fraction of the cost.
052 Recent work in EEG-based visual decoding (Li et al., 2024; Xiao et al., 2025; Zhang et al., 2025)
053 has demonstrated promising capabilities in both image retrieval and reconstruction, suggesting the
potential for more deployable applications.

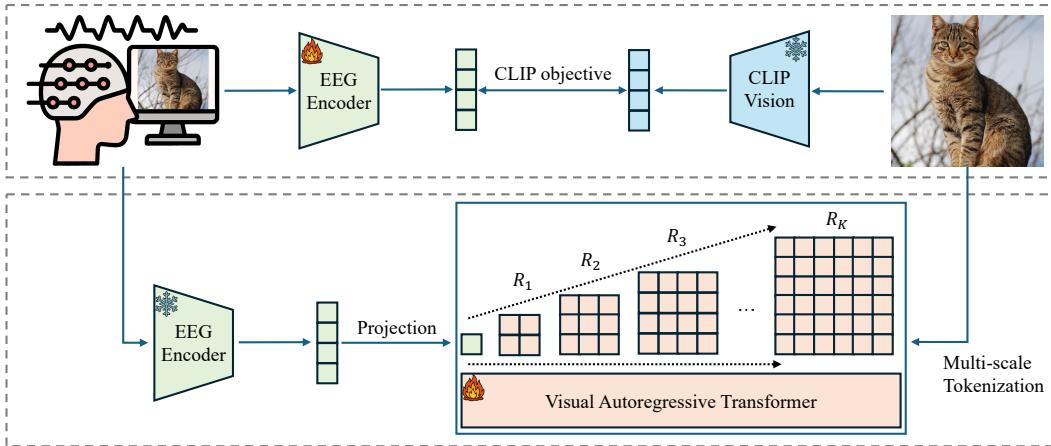


Figure 1: AVDE involves two training stages. **Stage 1:** A pre-trained EEG encoder is fine-tuned using contrastive learning to more effectively capture visual information embedded in EEG signals. This adaptation aims to provide a more informative initialization for the subsequent visual reconstruction process. **Stage 2:** A visual autoregressive transformer is trained using the next-scale prediction objective (Equation 7). Specifically, the model takes the sequence $([s], R_1, R_2, \dots, R_{K-1})$ as input and predicts the corresponding sequence $(R_1, R_2, R_3, \dots, R_K)$. Training is guided by a standard cross-entropy loss.

However, a fundamental challenge persists: how to effectively bridge the vast distributional gap between noisy EEG signals and structured visual content. This challenge manifests in three limitations of current approaches. First, these methods typically rely on diffusion models and complex, multi-stage adaptation processes based on unCLIP (Ramesh et al., 2022) to project EEG signals into compatible latent representations. The sequential nature of these pipelines inherently compounds errors across stages (Li & van der Schaar, 2023), degrading the fidelity of reconstructed images. Second, the EEG encoders are typically trained from scratch using a limited number of image-EEG pairs, which raises concerns about their capability to capture the intricate features in highly noisy EEG signals. Third, the computational demands of large-scale diffusion models (often exceeding 3B parameters) render these systems impractical for BCI applications where efficiency and responsiveness are crucial.

To address these limitations, we propose Autoregressive Visual Decoding from EEG signals (AVDE), a lightweight and efficient two-stage pipeline for EEG-to-image translation. Our approach makes two key innovations: First, rather than training EEG encoders from scratch, we leverage LaBraM (Jiang et al., 2024)—a model pre-trained on thousands of hours of diverse EEG data—and fine-tune it using contrastive learning to align EEG and image representations. This transfer learning approach substantially improves the extraction of meaningful features from noisy EEG signals. Second, we replace complex multi-stage diffusion processes with a streamlined autoregressive framework based on “next-scale prediction.” Our approach encodes images into multi-scale token maps using a pre-trained VQ-VAE (Tian et al., 2024), then trains a transformer to progressively predict increasingly detailed visual representations, starting from EEG embeddings as the coarsest representation. This approach ensures coherent generation while maintaining a direct relationship between EEG signals and visual outputs. Experiments on two datasets demonstrate that AVDE achieves state-of-the-art performance in both retrieval and reconstruction tasks while using only 10% of the parameters required by previous methods. Furthermore, visualization of the intermediate outputs shows that the generative process of AVDE reflects the hierarchical nature of human visual perception, underscoring the potential of autoregressive models as tools for exploring the dynamics of human visual cognition.

In summary, the main contributions are as follows:

- We introduce AVDE, a novel framework for EEG-based visual decoding that employs a hierarchical “next-scale prediction” strategy within an autoregressive transformer. This approach progressively constructs visual representations from coarse to fine details, mirroring

108 the hierarchical nature of both biological visual processing and computational vision sys-
 109 tems.
 110 • We demonstrate that transfer learning from pre-trained EEG model significantly improves
 111 visual decoding performance. By fine-tuning the LaBraM encoder (Jiang et al., 2024) with
 112 contrastive learning, we achieve more robust alignment between EEG and image represen-
 113 tation spaces compared to training EEG encoders from scratch.
 114 • We demonstrate through comprehensive experiments that AVDE achieves state-of-the-art
 115 performance in both image retrieval and reconstruction tasks on two datasets, while being
 116 more lightweight and computationally efficient than prior methods. Our approach reduces
 117 parameter count by approximately 90% compared to diffusion-based methods, making it
 118 more suitable for practical BCI applications.
 119

120 **2 METHOD**
 121

122 **2.1 EEG ENCODING WITH LABRAM**
 123

124 A critical challenge in EEG-based visual decoding is extracting meaningful features from the in-
 125 herently noisy signals. Rather than training encoders from scratch on limited EEG-image pairs, we
 126 build upon LaBraM (Jiang et al., 2024), a model pre-trained on over 2000 hours of diverse EEG data
 127 spanning multiple datasets and recording conditions.

128 The architecture processes input EEG data $X \in \mathbb{R}^{C \times T}$ (where C represents channels and T repre-
 129 sent time points) through the following encoding scheme:

130 1) **Temporal patching:** The input EEG signal is segmented in the temporal dimension with a non-
 131 overlapping window of length w , resulting in patches:
 132

$$\mathbf{x} = \{x_{c_j, k} \in \mathbb{R}^w \mid j = 1, 2, \dots, C, k = 1, 2, \dots, \lfloor \frac{T}{w} \rfloor\} \quad (1)$$

133 2) **Local feature extraction:** Each patch is processed by a temporal encoder comprising stacked
 134 convolutional blocks (1D convolution, group normalization, GELU activation) to capture fine-
 135 grained temporal patterns:
 136

$$\{e_{c_j, k} \in \mathbb{R}^d \mid j = 1, 2, \dots, C, k = 1, 2, \dots, \lfloor \frac{T}{w} \rfloor\} \quad (2)$$

137 where d is the embedding dimension.
 138

139 3) **Spatiotemporal contextualization:** To incorporate both temporal and spatial context into the
 140 model, we set up two sets of trainable positional embeddings: a temporal embedding set $TE =$
 141 $\{te_k \mid k = 1, 2, \dots, \lfloor \frac{T}{w} \rfloor\}$ and a spatial embedding set $SE = \{se_j \mid j = 1, 2, \dots, C\}$. The final
 142 patch representation is obtained by summing the corresponding temporal and spatial embeddings
 143 with the encoder output:
 144

$$\{e_{c_j, k} + te_k + se_j \mid j = 1, 2, \dots, C, k = 1, 2, \dots, \lfloor \frac{T}{w} \rfloor\} \quad (3)$$

145 4) **Global integration :** The enriched patch embeddings are processed by a Transformer encoder
 146 (Vaswani et al., 2017) that models dependencies across both time and channels, effectively integrat-
 147 ing information from the entire EEG epoch.
 148

149 **2.2 REPRESENTATION ALIGNMENT THROUGH CONTRASTIVE LEARNING**
 150

151 While pre-training provides a strong foundation for EEG feature extraction, the LaBraM model
 152 was primarily trained on clinical data (Obeid & Picone, 2016) rather than EEG responses to visual
 153 stimuli. To adapt the model for visual decoding, we fine-tune it through contrastive learning, which
 154 creates alignment between EEG and image representation spaces.
 155

156 Given paired EEG-image data ($\mathbf{X} \in \mathbb{R}^{B \times C \times T}, \mathbf{I} \in \mathbb{R}^{B \times H \times W}$), we encode EEG signals us-
 157 ing the LaBraM model and images using a frozen CLIP (Radford et al., 2021) encoder, produc-
 158 ing embeddings $\mathbf{e}, \mathbf{z} \in \mathbb{R}^{B \times d}$. We then optimize a bidirectional contrastive objective that max-
 159 imizes agreement between corresponding EEG-image pairs while minimizing similarity between
 160 non-corresponding pairs:
 161

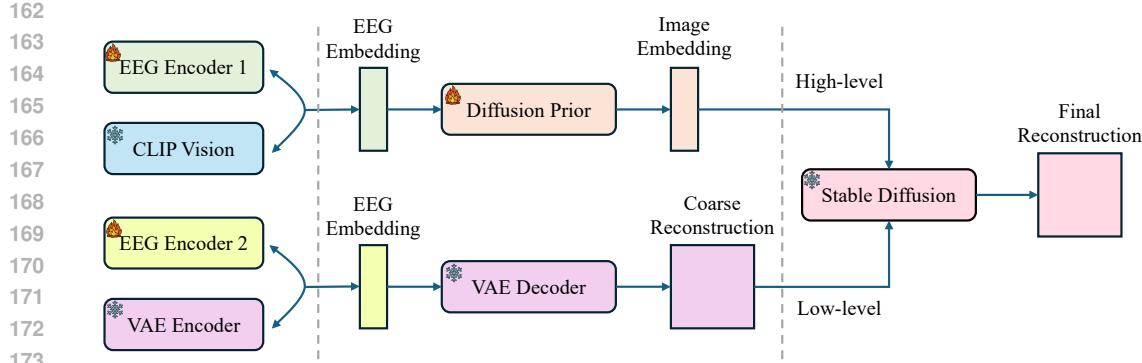


Figure 2: A typical unCLIP framework employed in previous EEG-based visual decoding works Li et al. (2024); Zhang et al. (2025); Xiao et al. (2025); Scotti et al. (2023). Despite its flexibility, the framework comprises multiple stages (five in this case), each introducing potential sources of error that can accumulate and degrade overall performance. Furthermore, the computational and memory demands of its components present significant challenges for practical implementation in BCIs.

$$\mathcal{L}_{CLIP} = -\frac{1}{B} \sum_{i=1}^B \left(\log \frac{\exp(s(e_i, z_i)/\tau)}{\sum_{j=1}^B \exp(s(e_i, z_j)/\tau)} + \log \frac{\exp(s(e_i, z_i)/\tau)}{\sum_{k=1}^B \exp(s(e_k, z_i)/\tau)} \right) \quad (4)$$

where s denotes cosine similarity and τ is a learned temperature parameter that controls the sharpness of the distribution. This objective effectively creates a shared embedding space where EEG signals are mapped near their corresponding image representations and away from unrelated ones.

To further strengthen the alignment, we incorporate a direct regression objective following practices established in Benchetrit et al. (2023) and Li et al. (2024):

$$\mathcal{L}_{Combined} = \lambda \mathcal{L}_{CLIP} + (1 - \lambda) \mathcal{L}_{MSE} \quad (5)$$

where \mathcal{L}_{MSE} is the mean squared error between normalized EEG and image embeddings, and λ (set to 0.8 in our experiments) balances the two objectives. This dual-objective approach provides more stable training dynamics and improves the precision of the EEG-to-image mapping by combining the structural alignment properties of contrastive learning with the point-wise precision of direct regression.

2.3 AUTOREGRESSIVE EEG-TO-IMAGE GENERATION

With aligned EEG representations in hand, we turn to the challenge of generating corresponding images. Rather than employing complex diffusion-based pipelines, we adopt a hierarchical autoregressive approach inspired by VAR (Tian et al., 2024). This framework enables direct, progressive image generation from EEG embeddings through a coarse-to-fine refinement process.

The architecture consists of two key components:

1) **Multi-scale image tokenization:** A pre-trained VQ-VAE tokenizes images into a hierarchy of discrete representations at multiple resolutions. Given an image I , the tokenizer produces a feature map $F \in \mathbb{R}^{h \times w \times d}$ that is quantized into K multi-scale residual maps (R_1, R_2, \dots, R_K) , where each R_k has resolution $h_k \times w_k$ that progressively increases with k .

These residual maps can be combined to progressively reconstruct the full-resolution feature map:

$$F_k = \sum_{i=1}^k \text{up}(R_i, (h, w)), \quad (6)$$

where $\text{up}(\cdot)$ denotes bilinear upsampling, and F_k represents the accumulated feature map after incorporating the first k residuals. This formulation allows the image to be constructed incrementally, from coarse structures to fine details.

216 2) **Next-scale prediction transformer:** A decoder-only transformer is trained to autoregressively
 217 predict these residual maps from EEG embeddings e . The model learns to generate increasingly
 218 detailed visual representations according to:

$$219 \quad 220 \quad 221 \quad 222 \quad 223 \quad 224 \quad 225 \quad 226 \quad 227 \quad 228 \quad 229 \quad 230 \quad 231 \quad 232 \quad 233 \quad 234 \quad 235 \quad 236 \quad 237 \quad 238 \quad 239 \quad 240 \quad 241 \quad 242 \quad 243 \quad 244 \quad 245 \quad 246 \quad 247 \quad 248 \quad 249 \quad 250 \quad 251 \quad 252 \quad 253 \quad 254 \quad 255 \quad 256 \quad 257 \quad 258 \quad 259 \quad 260 \quad 261 \quad 262 \quad 263 \quad 264 \quad 265 \quad 266 \quad 267 \quad 268 \quad 269$$

$$p(R_1, \dots, R_K) = \prod_{k=1}^K p(R_k | R_1, \dots, R_{k-1}, e), \quad (7)$$

where the sequence (R_1, \dots, R_{k-1}, e) provides the context for predicting the next-scale residual R_k .

This formulation is particularly appropriate for neural decoding because it mirrors theories of hierarchical visual processing in the brain, where perception progresses from coarse features to increasingly fine details. The EEG embedding e serves as the initial neural representation of the perceived image, and the transformer progressively elaborates this representation across multiple scales.

In practice, as shown in Fig. 1, the EEG embedding $e \in \mathbb{R}^d$ is first projected to the transformer’s hidden dimension h to create a special token $[s]$, which initiates the generation process. For each subsequent scale $k > 1$, the model processes the appropriately downsampled version of the previous cumulative feature map:

$$\tilde{F}_{k-1} = \text{down}(F_{k-1}, (h_k, w_k)), \quad (8)$$

where $\text{down}(\cdot)$ represents bilinear downsampling to match the target resolution (h_k, w_k) of the current scale.

During training, we employ a block-wise causal attention mask to ensure the model only attends to the appropriate context when predicting each scale. During inference, the process begins with the EEG embedding and autoregressively generates each scale until reaching the final resolution, at which point the multi-scale VQ-VAE decoder transforms the predicted feature map \tilde{F}_K into a complete image.

3 EXPERIMENTS

3.1 EXPERIMENTAL SETUP

We primarily evaluate our method on the THINGS-EEG dataset (Grootswagers et al., 2022), which serves as a widely adopted benchmark for EEG-based visual decoding. To further verify the versatility of AVDE, we additionally conduct experiments on the EEG-ImageNet dataset (Zhu et al., 2024), with results reported in Appendix C.

Dataset Overview. The THINGS-EEG dataset (Grootswagers et al., 2022) contains EEG recordings from 10 participants collected under a rapid serial visual presentation (RSVP) paradigm. The training set consists of 1,654 object concepts, each associated with 10 images presented four times, yielding a total of 66,160 EEG trials. The test set includes 200 distinct concepts, each represented by a single image repeated 80 times, resulting in 16,000 EEG trials. To mitigate habituation effects, both training and test images are presented in a pseudorandom order. Each image is displayed for 100 milliseconds, followed by a 100-millisecond blank screen to reduce blink-related and other artifacts. EEG signals were recorded from 63 channels, band-pass filtered between 0.1 Hz and 100 Hz, and sampled at 1,000 Hz.

Data Preprocessing. Following the practice in Song et al. (2023) and Li et al. (2024), we segment the EEG data into epochs spanning 0 to 1,000 ms relative to stimulus onset and apply baseline correction using the mean signal from the 200 ms pre-stimulus interval. All electrodes are preserved, and the data are downsampled to 200 Hz. Given that EEG amplitudes typically range from -0.1 mV to 0.1 mV, we normalize the signals by scaling them with respect to 0.1 mV, resulting in values primarily distributed between -1 and 1 . For the test set, EEG responses corresponding to each image are averaged across repetitions to improve the signal-to-noise ratio.

Implementation Details. We initialize the EEG encoder and the visual autoregressive (VAR) transformer with the pre-trained weights provided in the official GitHub repositories of LaBraM (Jiang et al., 2024) and VAR (Tian et al., 2024), respectively. The EEG encoder is trained using

270 Table 1: Overall accuracy of 200-way zero-shot retrieval under both within-subject and cross-subject
 271 settings. Each cell presents the Top-1 accuracy on the first line and the Top-5 accuracy on the second
 272 line. Results are averaged over five different random seeds; corresponding standard deviation values
 273 are presented in Table 12. For each subject, the highest accuracy values are indicated in bold.

Method	Sub-01	Sub-02	Sub-03	Sub-04	Sub-05	Sub-06	Sub-07	Sub-08	Sub-09	Sub-10	Ave
Within-subject: train and test on one subject											
EEGNetV4 (Lawhern et al., 2018)	0.144 0.391	0.159 0.398	0.202 0.432	0.224 0.517	0.132 0.289	0.129 0.402	0.198 0.467	0.246 0.549	0.184 0.419	0.237 0.543	0.186 0.441
EEGConformer (Song et al., 2022)	0.095 0.261	0.108 0.274	0.142 0.306	0.155 0.371	0.088 0.198	0.081 0.280	0.122 0.318	0.164 0.405	0.109 0.298	0.151 0.392	0.122 0.310
NICE (Song et al., 2023)	0.201 0.479	0.192 0.369	0.212 0.538	0.224 0.504	0.144 0.316	0.261 0.563	0.269 0.557	0.382 0.674	0.234 0.532	0.298 0.586	0.242 0.512
ATM (Li et al., 2024)	0.232 0.512	0.188 0.432	0.273 0.570	0.280 0.541	0.168 0.395	0.280 0.592	0.268 0.537	0.393 0.715	0.245 0.512	0.372 0.677	0.269 0.548
AVDE (Ours)	0.250 0.552	0.241 0.510	0.275 0.586	0.298 0.547	0.254 0.503	0.335 0.603	0.274 0.552	0.417 0.713	0.261 0.521	0.395 0.730	0.300 0.582
Cross-subject: leave one subject out for test											
EEGNetV4	0.086 0.232	0.082 0.226	0.073 0.171	0.113 0.257	0.092 0.217	0.101 0.224	0.056 0.182	0.084 0.231	0.074 0.196	0.124 0.305	0.089 0.224
EEGConformer	0.069 0.197	0.066 0.193	0.058 0.146	0.090 0.217	0.074 0.185	0.081 0.191	0.045 0.156	0.067 0.198	0.059 0.167	0.099 0.260	0.071 0.191
NICE	0.103 0.286	0.100 0.257	0.086 0.206	0.127 0.323	0.091 0.183	0.146 0.341	0.102 0.268	0.112 0.239	0.098 0.242	0.169 0.386	0.113 0.273
ATM	0.121 0.296	0.128 0.302	0.082 0.224	0.127 0.293	0.094 0.249	0.107 0.259	0.083 0.257	0.122 0.296	0.096 0.247	0.171 0.381	0.115 0.280
AVDE (Ours)	0.141 0.322	0.170 0.384	0.091 0.218	0.152 0.325	0.125 0.324	0.173 0.386	0.074 0.204	0.185 0.401	0.132 0.336	0.180 0.393	0.143 0.329

295 the AdamW optimizer with an initial learning rate of 2e-3, a weight decay of 0.05, and a minimum
 296 learning rate of 1e-5. The batch size is set to 1024. For the VAR transformer, we configure the model
 297 with a depth of 16 and train it using the AdamW optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.95$, a base learning
 298 rate of 2e-5, a weight decay of 0.05, a global batch size of 512, and 50 training epochs. Additional
 299 hyperparameter details are provided in Appendix B. During generation, we employ classifier-free
 300 guidance (CFG) with a ratio of 4.0 and apply top-k sampling with $k = 900$. All the experiments are
 301 conducted on Linux servers equipped with four NVIDIA A100 (40G) GPUs and Python 3.10.16 +
 302 PyTorch 2.5.1 + CUDA 12.4 environment.

303 **Evaluation.** We assess the effectiveness of AVDE on both image retrieval and reconstruction tasks.
 304 For the retrieval task, we compute the cosine similarity between the EEG embeddings generated by
 305 the EEG encoder and the CLIP image embeddings of 200 test concepts. Retrieval performance is
 306 evaluated based on the probability that the ground truth concept appears among the top-K candidates
 307 ($K = 1$ or 5). For the reconstruction task, we adopt standard evaluation metrics following prior work
 308 (Scotti et al., 2023; Li et al., 2024) to quantify the similarity between reconstructed and ground truth
 309 visual stimuli: (1) PixCorr – pixel-wise correlation; (2) SSIM – Structural Similarity Index Measure;
 310 (3) SwAV – average correlation distance computed from SwAV-ResNet50 (Caron et al., 2020) fea-
 311 tures; and (4) Two-way identification using pretrained neural networks (AlexNet (Krizhevsky et al.,
 312 2012) layers 2 and 5, Inception (Szegedy et al., 2015), and CLIP). Two-way identification is treated
 313 as a bidirectional retrieval task, as described in Ozcelik & VanRullen (2023).

315 3.2 RETRIEVAL PERFORMANCE

317 Table 1 presents a quantitative evaluation of EEG-based image retrieval performance, comparing
 318 our proposed method, AVDE, with several baseline approaches. Remarkably, AVDE achieves a
 319 top-1 accuracy of 0.300 and a top-5 accuracy of 0.582 in the zero-shot EEG-to-image retrieval task
 320 under the within-subject setting. Under the more challenging cross-subject setting, it attains a top-1
 321 accuracy of 0.143 and a top-5 accuracy of 0.329. These results represent a substantial improvement
 322 over existing state-of-the-art methods, highlighting the effectiveness of our approach.

323 The strong performance of AVDE underscores the utility of the LaBraM-based EEG encoder, which
 324 benefits significantly from large-scale pre-training. This pre-training enables the encoder to gen-

324
 325
 326
 327
 328
 329
 330
 331
 332
 333
 334
 335
 336
 337
 338
 339
 340
 341
 342
 343
 344
 345
 346
 347
 348
 349
 350
 351
 352
 353
 354
 355
 356
 357
 358
 359
 360
 361
 362
 363
 364
 365
 366
 367
 368
 369
 370
 371
 372
 373
 374
 375
 376
 377
 378
 379
 380
 381
 382
 383
 384
 385
 386
 387
 388
 389
 390
 391
 392
 393
 394
 395
 396
 397
 398
 399
 400
 401
 402
 403
 404
 405
 406
 407
 408
 409
 410
 411
 412
 413
 414
 415
 416
 417
 418
 419
 420
 421
 422
 423
 424
 425
 426
 427
 428
 429
 430
 431
 432
 433
 434
 435
 436
 437
 438
 439
 440
 441
 442
 443
 444
 445
 446
 447
 448
 449
 450
 451
 452
 453
 454
 455
 456
 457
 458
 459
 460
 461
 462
 463
 464
 465
 466
 467
 468
 469
 470
 471
 472
 473
 474
 475
 476
 477
 478
 479
 480
 481
 482
 483
 484
 485
 486
 487
 488
 489
 490
 491
 492
 493
 494
 495
 496
 497
 498
 499
 500
 501
 502
 503
 504
 505
 506
 507
 508
 509
 510
 511
 512
 513
 514
 515
 516
 517
 518
 519
 520
 521
 522
 523
 524
 525
 526
 527
 528
 529
 530
 531
 532
 533
 534
 535
 536
 537
 538
 539
 540
 541
 542
 543
 544
 545
 546
 547
 548
 549
 550
 551
 552
 553
 554
 555
 556
 557
 558
 559
 560
 561
 562
 563
 564
 565
 566
 567
 568
 569
 570
 571
 572
 573
 574
 575
 576
 577
 578
 579
 580
 581
 582
 583
 584
 585
 586
 587
 588
 589
 590
 591
 592
 593
 594
 595
 596
 597
 598
 599
 600
 601
 602
 603
 604
 605
 606
 607
 608
 609
 610
 611
 612
 613
 614
 615
 616
 617
 618
 619
 620
 621
 622
 623
 624
 625
 626
 627
 628
 629
 630
 631
 632
 633
 634
 635
 636
 637
 638
 639
 640
 641
 642
 643
 644
 645
 646
 647
 648
 649
 650
 651
 652
 653
 654
 655
 656
 657
 658
 659
 660
 661
 662
 663
 664
 665
 666
 667
 668
 669
 670
 671
 672
 673
 674
 675
 676
 677
 678
 679
 680
 681
 682
 683
 684
 685
 686
 687
 688
 689
 690
 691
 692
 693
 694
 695
 696
 697
 698
 699
 700
 701
 702
 703
 704
 705
 706
 707
 708
 709
 710
 711
 712
 713
 714
 715
 716
 717
 718
 719
 720
 721
 722
 723
 724
 725
 726
 727
 728
 729
 730
 731
 732
 733
 734
 735
 736
 737
 738
 739
 740
 741
 742
 743
 744
 745
 746
 747
 748
 749
 750
 751
 752
 753
 754
 755
 756
 757
 758
 759
 760
 761
 762
 763
 764
 765
 766
 767
 768
 769
 770
 771
 772
 773
 774
 775
 776
 777
 778
 779
 780
 781
 782
 783
 784
 785
 786
 787
 788
 789
 790
 791
 792
 793
 794
 795
 796
 797
 798
 799
 800
 801
 802
 803
 804
 805
 806
 807
 808
 809
 810
 811
 812
 813
 814
 815
 816
 817
 818
 819
 820
 821
 822
 823
 824
 825
 826
 827
 828
 829
 830
 831
 832
 833
 834
 835
 836
 837
 838
 839
 840
 841
 842
 843
 844
 845
 846
 847
 848
 849
 850
 851
 852
 853
 854
 855
 856
 857
 858
 859
 860
 861
 862
 863
 864
 865
 866
 867
 868
 869
 870
 871
 872
 873
 874
 875
 876
 877
 878
 879
 880
 881
 882
 883
 884
 885
 886
 887
 888
 889
 890
 891
 892
 893
 894
 895
 896
 897
 898
 899
 900
 901
 902
 903
 904
 905
 906
 907
 908
 909
 910
 911
 912
 913
 914
 915
 916
 917
 918
 919
 920
 921
 922
 923
 924
 925
 926
 927
 928
 929
 930
 931
 932
 933
 934
 935
 936
 937
 938
 939
 940
 941
 942
 943
 944
 945
 946
 947
 948
 949
 950
 951
 952
 953
 954
 955
 956
 957
 958
 959
 960
 961
 962
 963
 964
 965
 966
 967
 968
 969
 970
 971
 972
 973
 974
 975
 976
 977
 978
 979
 980
 981
 982
 983
 984
 985
 986
 987
 988
 989
 990
 991
 992
 993
 994
 995
 996
 997
 998
 999
 1000
 1001
 1002
 1003
 1004
 1005
 1006
 1007
 1008
 1009
 1010
 1011
 1012
 1013
 1014
 1015
 1016
 1017
 1018
 1019
 1020
 1021
 1022
 1023
 1024
 1025
 1026
 1027
 1028
 1029
 1030
 1031
 1032
 1033
 1034
 1035
 1036
 1037
 1038
 1039
 1040
 1041
 1042
 1043
 1044
 1045
 1046
 1047
 1048
 1049
 1050
 1051
 1052
 1053
 1054
 1055
 1056
 1057
 1058
 1059
 1060
 1061
 1062
 1063
 1064
 1065
 1066
 1067
 1068
 1069
 1070
 1071
 1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079
 1080
 1081
 1082
 1083
 1084
 1085
 1086
 1087
 1088
 1089
 1090
 1091
 1092
 1093
 1094
 1095
 1096
 1097
 1098
 1099
 1100
 1101
 1102
 1103
 1104
 1105
 1106
 1107
 1108
 1109
 1110
 1111
 1112
 1113
 1114
 1115
 1116
 1117
 1118
 1119
 1120
 1121
 1122
 1123
 1124
 1125
 1126
 1127
 1128
 1129
 1130
 1131
 1132
 1133
 1134
 1135
 1136
 1137
 1138
 1139
 1140
 1141
 1142
 1143
 1144
 1145
 1146
 1147
 1148
 1149
 1150
 1151
 1152
 1153
 1154
 1155
 1156
 1157
 1158
 1159
 1160
 1161
 1162
 1163
 1164
 1165
 1166
 1167
 1168
 1169
 1170
 1171
 1172
 1173
 1174
 1175
 1176
 1177
 1178
 1179
 1180
 1181
 1182
 1183
 1184
 1185
 1186
 1187
 1188
 1189
 1190
 1191
 1192
 1193
 1194
 1195
 1196
 1197
 1198
 1199
 1200
 1201
 1202
 1203
 1204
 1205
 1206
 1207
 1208
 1209
 1210
 1211
 1212
 1213
 1214
 1215
 1216
 1217
 1218
 1219
 1220
 1221
 1222
 1223
 1224
 1225
 1226
 1227
 1228
 1229
 1230
 1231
 1232
 1233
 1234
 1235
 1236
 1237
 1238
 1239
 1240
 1241
 1242
 1243
 1244
 1245
 1246
 1247
 1248
 1249
 1250
 1251
 1252
 1253
 1254
 1255
 1256
 1257
 1258
 1259
 1260
 1261
 1262
 1263
 1264
 1265
 1266
 1267
 1268
 1269
 1270
 1271
 1272
 1273
 1274
 1275
 1276
 1277
 1278
 1279
 1280
 1281
 1282
 1283
 1284
 1285
 1286
 1287
 1288
 1289
 1290
 1291
 1292
 1293
 1294
 1295
 1296
 1297
 1298
 1299
 1300
 1301
 1302
 1303
 1304
 1305
 1306
 1307
 1308
 1309
 1310
 1311
 1312
 1313
 1314
 1315
 1316
 1317
 1318
 1319
 1320
 1321
 1322
 1323
 1324
 1325
 1326
 1327
 1328
 1329
 1330
 1331
 1332
 1333
 1334
 1335
 1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349
 1350
 1351
 1352
 1353
 1354
 1355
 1356
 1357
 1358
 1359
 1360
 1361
 1362
 1363
 1364
 1365
 1366
 1367
 1368
 1369
 1370
 1371
 1372
 1373
 1374
 1375
 1376
 1377
 1378
 1379
 1380
 1381
 1382
 1383
 1384
 1385
 1386
 1387
 1388
 1389
 1390
 1391
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403
 1404
 1405
 1406
 1407
 1408
 1409
 1410
 1411
 1412
 1413
 1414
 1415
 1416
 1417
 1418
 1419
 1420
 1421
 1422
 1423
 1424
 1425
 1426
 1427
 1428
 1429
 1430
 1431
 1432
 1433
 1434
 1435
 1436
 1437
 1438
 1439
 1440
 1441
 1442
 1443
 1444
 1445
 1446
 1447
 1448
 1449
 1450
 1451
 1452
 1453
 1454
 1455
 1456
 1457
 1458
 1459
 1460
 1461
 1462
 1463
 1464
 1465
 1466
 1467
 1468
 1469
 1470
 1471
 1472
 1473
 1474
 1475
 1476
 1477
 1478
 1479
 1480
 1481
 1482
 1483
 1484
 1485
 1486
 1487
 1488
 1489
 1490
 1491
 1492
 1493
 1494
 1495
 1496
 1497
 1498
 1499
 1500
 1501
 1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511
 1512
 1513
 1514
 1515
 1516
 1517
 1518
 1519
 1520
 1521
 1522
 1523
 1524
 1525
 1526
 1527
 1528
 1529
 1530
 1531
 1532
 1533
 1534
 1535
 1536
 1537
 1538
 1539
 1540
 1541
 1542
 1543
 1544
 1545
 1546
 1547
 1548
 1549
 1550
 1551
 1552
 1553
 1554
 1555
 1556
 1557
 1558
 1559
 1560
 1561
 1562
 1563
 1564
 1565
 1566
 1567
 1568
 1569
 1570
 1571
 1572
 1573
 1574
 1575
 1576
 1577
 1578
 1579
 1580
 1581
 1582
 1583
 1584
 1585
 1586
 1587
 1588
 1589
 1590
 1591
 1592
 1593
 1594
 1595
 1596
 1597
 1598
 1599
 1600
 1601
 1602
 1603
 1604
 1605
 1606
 1607
 1608
 1609
 1610
 1611
 1612
 1613
 1614
 1615
 1616
 1617
 1618
 1619
 1620
 1621
 1622
 1623
 1624
 1625
 1626
 1627
 1628
 1629
 1630
 1631
 1632
 1633
 1634
 1635
 1636
 1637
 1638
 1639
 1640
 1641
 1642
 1643
 1644
 1645
 1646
 1647
 1648
 1649
 1650
 1651
 1652
 1653
 1654
 1655
 1656
 1657
 1658
 1659
 1660
 1661
 1662
 1663
 1664
 1665
 1666
 1667
 1668
 1669
 1670
 1671
 1672
 1673
 1674
 1675
 1676
 1677
 1678
 1679
 1680
 1681
 1682
 1683
 1684
 1685
 1686
 1687
 1688
 1689
 1690
 1691
 1692
 1693
 1694
 1695
 1696
 1697
 1698
 1699
 1700
 1701
 1702
 1703
 1704
 1705
 1706
 1707
 1708
 1709
 1710
 1711
 1712
 1713
 1714
 1715
 1716
 1717
 1718
 1719
 1720
 1721
 1722
 1723
 1724
 1725
 1726
 1727
 1728
 1729
 1730
 1731
 1732
 1733
 1734
 1735
 1736
 1737
 1738
 1739
 1740
 1741
 1742
 1743
 1744
 1745
 1746
 1747
 1748
 1749
 1750
 1751
 1752
 1753
 1754
 1755
 1756
 1757
 1758
 1759
 1760
 1761
 1762
 1763
 1764
 1765
 1766
 1767
 1768
 1769
 1770
 1771
 1772
 1773
 1774
 1775
 1776
 1777
 1778
 1779
 1780
 1781
 1782
 1783
 1784
 1785
 1786
 1787
 1788
 1789
 1790
 1791
 1792
 1793
 1794
 1795
 1796
 1797
 1798
 1799
 1800
 1801
 1802
 1803
 1804
 1805
 1806
 1807
 1808
 1809
 1810
 1811
 1812
 1813
 1814
 1815
 1816
 1817
 1818
 1819
 1820
 1821
 1822
 1823<br

378 diffusion model (SDXL (Podell et al., 2023)) for image generation, we select Li et al. (2024) as
 379 a representative due to its widespread adoption and open-source availability. To comprehensively
 380 assess both computational and spatial efficiency, we evaluate the following metrics: (1) FLOPs —
 381 the number of floating-point operations; (2) Inference time — the GPU time required for generation;
 382 and (3) Memory usage — the peak GPU memory usage during inference. All metrics are measured
 383 using PyTorch’s built-in profiler on a single NVIDIA A100 GPU. The batch size is set to 1, cor-
 384 responding to the resource cost of generating a single image. As summarized in Table 3, AVDE
 385 achieves faster image generation and lower memory consumption compared to prior state-of-the-art
 386 methods, demonstrating its superior suitability for practical applications.

387 Table 3: Comparison with state-of-the-art method on inference efficiency. All metrics are evaluated
 388 using PyTorch’s built-in profiler on a single NVIDIA A100 GPU. The batch size is fixed at 1, and
 389 results are averaged over 200 runs to ensure stability and reliability.

Method	Params (M)	Steps	FLOPs (G)	Inference Time (ms)	Memory Usage (MB)
Li et al. (2024)	3818.1	4	8738.6	310.4	4826.73
AVDE (Ours)	425.3	10	1350.5	91.2	1809.63

396 3.5 ABLATION STUDY

397 To evaluate the contribution of each core component in AVDE, namely the pre-trained EEG encoder
 398 and the autoregressive generative framework, we perform the following experiments: (1) Encoder
 399 substitution (ATM/EEGNet/NICE + VAR): These experiments replace the LaBraM encoder with
 400 other widely used EEG encoders (2) Generative framework substitution (LaBraM + Li et al., unCLIP
 401 baseline): This setting replaces the VAR generative framework with a standard unCLIP pipeline (Li
 402 et al., 2024). (3) Model substitution with diffusion models (LaBraM + LDM-4 (Rombach et al.,
 403 2022) / DiT-XL (Peebles & Xie, 2023)): These experiments replace the VAR model with slightly
 404 larger diffusion models trained under the same conditions. As shown in Table 4, performance de-
 405 grades when the EEG encoder is replaced, underscoring the value of high-quality embeddings from
 406 the pre-trained encoder for accurate visual reconstruction. Similarly, replacing the autoregressive
 407 framework results in a substantial drop in performance, suggesting that our overall training strategy
 408 more effectively aligns the distributional characteristics of EEG signals with those of natural images.

409
 410 Table 4: Impact of using different EEG encoders or generative framework on the reconstruction
 411 performance. The results are averaged over all subjects.

Method	Low-level			High-level			
	PixCorr \uparrow	SSIM \uparrow	AlexNet(2) \uparrow	AlexNet(5) \uparrow	Inception \uparrow	CLIP \uparrow	SwAV \downarrow
LaBraM+VAR	0.147	0.366	0.766	0.835	0.724	0.747	0.586
ATM+VAR	0.141	0.351	0.752	0.821	0.711	0.731	0.601
EEGNet+VAR	0.132	0.323	0.733	0.803	0.687	0.712	0.627
NICE+VAR	0.136	0.341	0.742	0.812	0.701	0.719	0.613
LaBraM+Li et al. (2024)	0.138	0.346	0.746	0.817	0.707	0.726	0.606
LaBraM+LDM-4	0.139	0.343	0.750	0.825	0.713	0.731	0.609
LaBraM+DiT-XL/2	0.143	0.354	0.761	0.829	0.715	0.735	0.594

423 424 3.6 ANALYSIS OF INTERMEDIATE OUTPUTS

425 Given that the “next-scale prediction” strategy employed in AVDE constitutes a progressive gener-
 426 ative process, we examine how the model incrementally extracts and interprets visual information
 427 from EEG signals throughout this procedure. To this end, we visualize all intermediate reconstruc-
 428 tions by accumulating the feature maps at each scale and decoding them into images using the
 429 decoder of a pre-trained multi-scale VQ-VAE. Formally, the set of cumulative feature maps is de-
 430 noted as $\{F_k \mid k = 1, 2, \dots, K\}$, where K represents the total number of scales, and each F_k is
 431 computed as described in Equation 6.

As shown in Figure 4, the generative process in AVDE exhibits notable parallels to the hierarchical organization of human visual perception. In the early stages of generation, the model produces coarse features—mirroring the role of the retina and primary visual cortex (V1), which primarily encode low-level visual attributes such as edges and color gradients (Tong, 2003; Tootell et al., 1998). As the process continues, mid-level features begin to emerge, resembling the functional role of V2 and V4 in integrating contours and object-level structures (Hegdé & Van Essen, 2007). In the final stages, the model reconstructs semantically rich and coherent imagery, analogous to the activity in higher-order visual regions such as the inferotemporal cortex, where holistic object representations are formed (Tanaka, 1996).

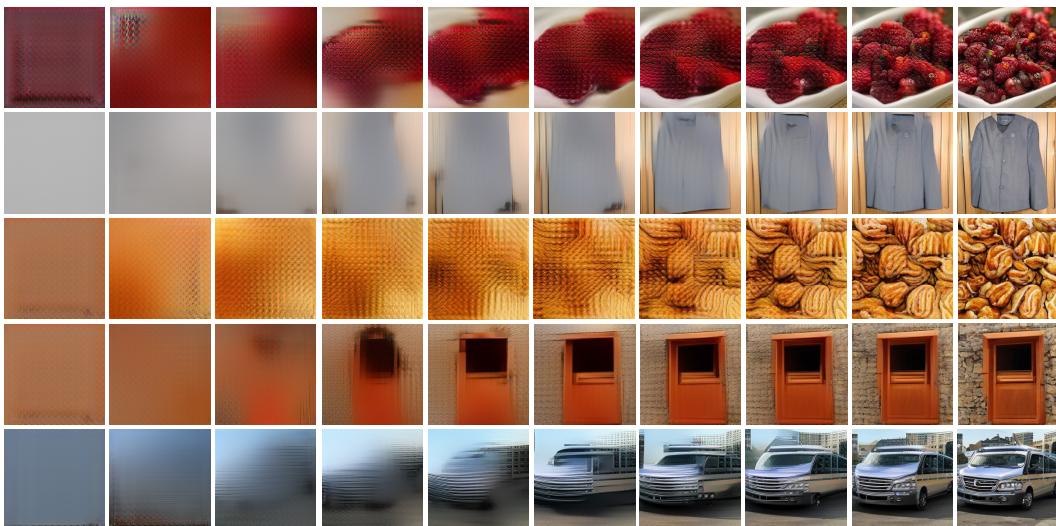


Figure 4: Intermediate reconstructions generated by AVDE across 10 progressive scales. Each row corresponds to a distinct EEG-evoked reconstruction instance, and each column represents the cumulative output up to a given scale. This process reflects the hierarchical nature of human visual perception, drawing parallels to the function of successive cortical visual areas (e.g., V1, V2/V4, and IT).

To examine the contribution of each intermediate scale to the generative process, we quantify the correlations between intermediate image features and EEG features derived from different brain regions. As shown in Figure 5a, EEG electrodes are grouped into five regions: frontal, temporal, central, parietal, and occipital. For each region, mean channel embeddings are computed using the EEG encoder, while for each intermediate scale, image embeddings are obtained from the CLIP image encoder. The cosine similarity between each region–scale pair is then calculated to assess their correspondence.

Since the generative process is cumulative, we compute the stepwise increase at each scale to capture the incremental information contributed by that scale. As shown in Figure 5c, the step increase for occipital regions peaks at early scales and gradually diminishes thereafter. In contrast, the temporal and parietal regions exhibit relatively sustained step increases across early and middle scales, followed by a decline in later scales. The frontal and central regions, however, show low step increases initially, which progressively rise and peak at late scales. These results suggest that the intermediate scales reflect the functional roles of different brain regions during visual processing.

4 CONCLUSION

In this work, we presented AVDE, a novel autoregressive framework for visual decoding from EEG signals that addresses key limitations of existing approaches in terms of complexity, efficiency, and performance. By leveraging pre-trained EEG representations via LaBraM and replacing multi-stage diffusion pipelines with a streamlined autoregressive process, AVDE enables accurate and coherent reconstruction of visual content from noisy EEG data. Experiments on two datasets demonstrate that AVDE outperforms state-of-the-art approaches in both retrieval and reconstruction tasks while

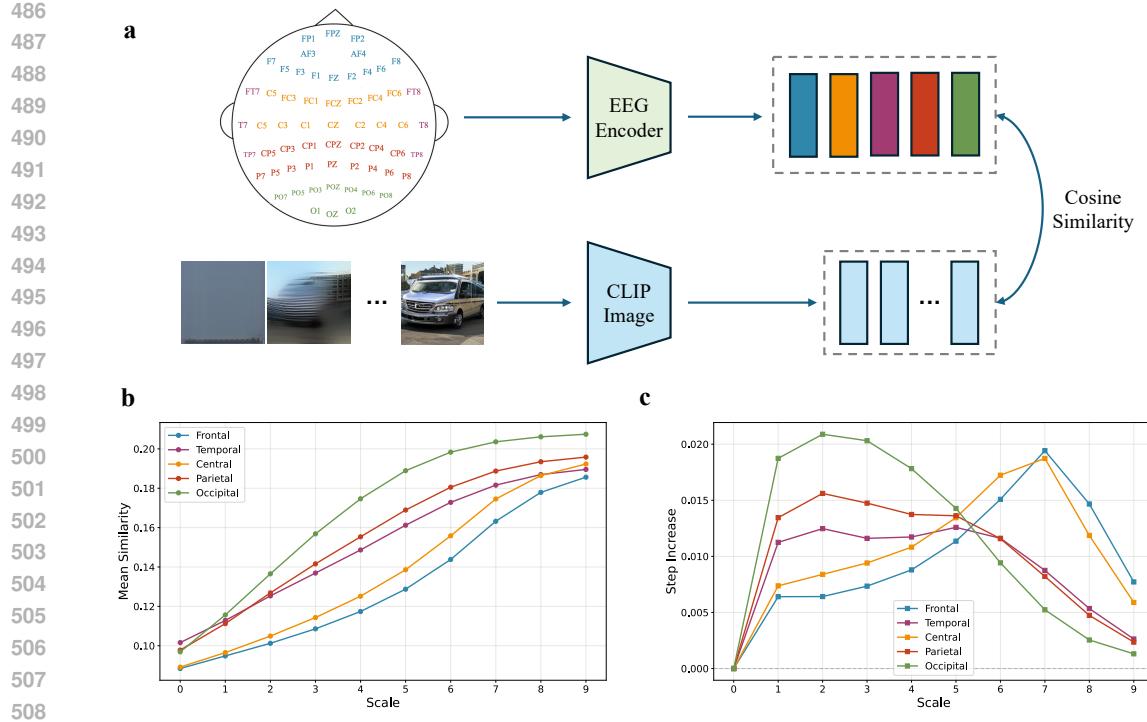


Figure 5: Analysis of similarities between intermediate scales and brain regions. (a) The mean channel embeddings from five brain regions are compared with the intermediate image embeddings. Cosine similarity is used as the measure. (b) Since the generative process is cumulative, the similarities generally increase as more scales are involved. (c) Stepwise increase captures the incremental information contributed by each scale. The step increase for occipital regions peaks at early scales and gradually diminishes thereafter. The temporal and parietal regions exhibit relatively sustained step increases across early and middle scales, followed by a decline in later scales. The frontal and central regions show low step increases initially, which progressively rise and peak at late scales.

requiring only a fraction of their computational resources, making it well-suited for practical BCI applications. Moreover, the hierarchical structure of AVDE's generative process mirrors the structure of human visual perception, highlighting its potential as a computational tool for investigating the mechanisms of visual cognition.

REFERENCES

Yohann Benchirrit, Hubert Banville, and Jean-Rémi King. Brain decoding: toward real-time reconstruction of visual perception. *arXiv preprint arXiv:2310.19812*, 2023.

Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in neural information processing systems*, 33:9912–9924, 2020.

Radoslaw Martin Cichy and Dimitrios Pantazis. Multivariate pattern analysis of meg and eeg: A comparison of representational structure in time and space. *NeuroImage*, 158:441–454, 2017.

Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12873–12883, 2021.

Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024.

540 Tao Fang, Qian Zheng, and Gang Pan. Alleviating the semantic gap for generalized fmri-to-image
 541 reconstruction. *Advances in Neural Information Processing Systems*, 36:15096–15107, 2023.
 542

543 Tijl Grootswagers, Ivy Zhou, Amanda K Robinson, Martin N Hebart, and Thomas A Carlson. Hu-
 544 man eeg recordings for 1,854 concepts presented in rapid serial visual presentation streams. *Sci-
 545 entific Data*, 9(1):3, 2022.

546 Martin N Hebart, Oliver Contier, Lina Teichmann, Adam H Rockter, Charles Y Zheng, Alexis Kid-
 547 der, Anna Corriveau, Maryam Vaziri-Pashkam, and Chris I Baker. Things-data, a multimodal
 548 collection of large-scale datasets for investigating object representations in human brain and be-
 549 havior. *Elife*, 12:e82580, 2023.

550 Jay Hegde and David C Van Essen. A comparative study of shape representation in macaque visual
 551 areas v2 and v4. *Cerebral cortex*, 17(5):1100–1116, 2007.
 552

553 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in
 554 neural information processing systems*, 33:6840–6851, 2020.
 555

556 Wei-Bang Jiang, Li-Ming Zhao, and Bao-Liang Lu. Large brain model for learning generic repre-
 557 sentations with tremendous eeg data in bci. *arXiv preprint arXiv:2405.18765*, 2024.
 558

559 Kendrick N Kay, Thomas Naselaris, Ryan J Prenger, and Jack L Gallant. Identifying natural images
 560 from human brain activity. *Nature*, 452(7185):352–355, 2008.
 561

562 Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convo-
 563 lutional neural networks. *Advances in neural information processing systems*, 25, 2012.
 564

565 Vernon J Lawhern, Amelia J Solon, Nicholas R Waytowich, Stephen M Gordon, Chou P Hung, and
 566 Brent J Lance. Eegnet: a compact convolutional neural network for eeg-based brain–computer
 567 interfaces. *Journal of neural engineering*, 15(5):056013, 2018.
 568

569 Doyup Lee, Chiheon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Autoregressive image
 570 generation using residual quantization. In *Proceedings of the IEEE/CVF Conference on Computer
 571 Vision and Pattern Recognition*, pp. 11523–11532, 2022.
 572

573 Dongyang Li, Chen Wei, Shiying Li, Jiachen Zou, Haoyang Qin, and Quanying Liu. Visual decoding
 574 and reconstruction via eeg embeddings with guided diffusion. *arXiv preprint arXiv:2403.07721*,
 575 2024.
 576

577 Yangming Li and Mihaela van der Schaar. On error propagation of diffusion models. *arXiv preprint
 578 arXiv:2308.05021*, 2023.
 579

580 Yulong Liu, Yongqiang Ma, Wei Zhou, Guibo Zhu, and Nanning Zheng. Brainclip: Bridging brain
 581 and visual-linguistic representation via clip for generic natural visual stimulus decoding. *arXiv
 582 preprint arXiv:2302.12971*, 2023.
 583

584 Nikos K Logothetis. What we can do and what we cannot do with fmri. *Nature*, 453(7197):869–878,
 585 2008.
 586

587 Ravi S Menon and Seong-Gi Kim. Spatial and temporal limits in cognitive neuroimaging with fmri.
 588 *Trends in cognitive sciences*, 3(6):207–216, 1999.
 589

590 Yoichi Miyawaki, Hajime Uchida, Okito Yamashita, Masa-aki Sato, Yusuke Morito, Hiroki C Tan-
 591 abe, Norihiro Sadato, and Yukiyasu Kamitani. Visual image reconstruction from human brain
 592 activity using a combination of multiscale local image decoders. *Neuron*, 60(5):915–929, 2008.
 593

594 Thomas Naselaris, Ryan J Prenger, Kendrick N Kay, Michael Oliver, and Jack L Gallant. Bayesian
 595 reconstruction of natural images from human brain activity. *Neuron*, 63(6):902–915, 2009.
 596

597 Iyad Obeid and Joseph Picone. The temple university hospital eeg data corpus. *Frontiers in neuro-
 598 science*, 10:196, 2016.
 599

600 Furkan Ozcelik and Rufin VanRullen. Natural scene reconstruction from fmri signals using genera-
 601 tive latent diffusion. *Scientific Reports*, 13(1):15666, 2023.
 602

594 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of*
 595 *the IEEE/CVF international conference on computer vision*, pp. 4195–4205, 2023.
 596

597 Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe
 598 Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image
 599 synthesis. *arXiv preprint arXiv:2307.01952*, 2023.

600 Xuelin Qian, Yun Wang, Jingyang Huo, Jianfeng Feng, and Yanwei Fu. fmri-pte: A large-scale
 601 fmri pretrained transformer encoder for multi-subject brain activity decoding. *arXiv preprint*
 602 *arXiv:2311.00342*, 2023.
 603

604 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 605 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 606 models from natural language supervision. In *International conference on machine learning*, pp.
 607 8748–8763. PMLR, 2021.

608 Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-
 609 conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.
 610

611 Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with
 612 vq-vae-2. *Advances in neural information processing systems*, 32, 2019.

613 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 614 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF confer-
 615 ence on computer vision and pattern recognition*, pp. 10684–10695, 2022.

616

617 Paul Scotti, Atmadeep Banerjee, Jimmie Goode, Stepan Shabalin, Alex Nguyen, Aidan Dempster,
 618 Nathalie Verlinde, Elad Yundler, David Weisberg, Kenneth Norman, et al. Reconstructing the
 619 mind’s eye: fmri-to-image with contrastive learning and diffusion priors. *Advances in Neural
 620 Information Processing Systems*, 36:24705–24728, 2023.

621 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv
 622 preprint arXiv:2010.02502*, 2020.
 623

624 Yonghao Song, Qingqing Zheng, Bingchuan Liu, and Xiaorong Gao. Eeg conformer: Convolu-
 625 tional transformer for eeg decoding and visualization. *IEEE Transactions on Neural Systems and
 626 Rehabilitation Engineering*, 31:710–719, 2022.

627

628 Yonghao Song, Bingchuan Liu, Xiang Li, Nanlin Shi, Yijun Wang, and Xiaorong Gao. Decoding
 629 natural images from eeg for object recognition. *arXiv preprint arXiv:2308.13234*, 2023.

630

631 Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Du-
 632 mitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In
 633 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1–9, 2015.

634

635 Yu Takagi and Shinji Nishimoto. High-resolution image reconstruction with latent diffusion models
 636 from human brain activity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and
 Pattern Recognition*, pp. 14453–14463, 2023a.

637

638 Yu Takagi and Shinji Nishimoto. Improving visual image reconstruction from human brain activity
 639 using latent diffusion models via multiple decoded inputs. *arXiv preprint arXiv:2306.11536*,
 640 2023b.

641

642 Keiji Tanaka. Inferotemporal cortex and object vision. *Annual review of neuroscience*, 19(1):109–
 139, 1996.

643

644 Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling:
 645 Scalable image generation via next-scale prediction. *Advances in neural information processing
 646 systems*, 37:84839–84865, 2024.

647

Frank Tong. Primary visual cortex and visual awareness. *Nature reviews neuroscience*, 4(3):219–
 229, 2003.

648 Roger BH Tootell, Nouchine K Hadjikhani, Wim Vanduffel, Arthur K Liu, Janine D Mendola,
649 Martin I Sereno, and Anders M Dale. Functional analysis of primary visual cortex (v1) in humans.
650 *Proceedings of the National Academy of Sciences*, 95(3):811–817, 1998.

651

652 Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in
653 neural information processing systems*, 30, 2017.

654

655 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
656 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural informa-
657 tion processing systems*, 30, 2017.

658

659 Weihao Xia, Raoul De Charette, Cengiz Oztireli, and Jing-Hao Xue. Dream: Visual decoding
660 from reversing human visual system. In *Proceedings of the IEEE/CVF Winter Conference on
Applications of Computer Vision*, pp. 8226–8235, 2024.

661

662 Xin Xiao, Kaiwen Wei, Jiang Zhong, Xuekai Wei, and Jielu Yan. Eeg decoding and visual recon-
663 struction via 3d geometric with nonstationarity modelling. In *ICASSP 2025-2025 IEEE Interna-
664 tional Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5. IEEE, 2025.

665

666 Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan,
667 Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive models for content-
668 rich text-to-image generation. *arXiv preprint arXiv:2206.10789*, 2(3):5, 2022.

669

670 Kaifan Zhang, Lihuo He, Xin Jiang, Wen Lu, Di Wang, and Xinbo Gao. Cognitioncapturer: De-
671 coding visual stimuli from human eeg signal with multimodal information. In *Proceedings of the
AAAI Conference on Artificial Intelligence*, pp. 14486–14493, 2025.

672

673 Shuqi Zhu, Ziyi Ye, Qingyao Ai, and Yiqun Liu. Eeg-imagenet: An electroencephalogram
674 dataset and benchmarks with image visual stimuli of multi-granularity labels. *arXiv preprint
arXiv:2406.07151*, 2024.

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 A RELATED WORK
703
704705 A.1 VISUAL DECODING FROM NEURAL SIGNALS
706707 The field of visual decoding from neural signals has evolved substantially, progressing from early
708 pattern recognition techniques to increasingly sophisticated generative models. Foundational studies
709 (Kay et al., 2008; Naselaris et al., 2009) established key methodologies by employing Gabor
710 wavelet-based encoding models and Bayesian inference to identify stimuli from fMRI activity pat-
711 terns. These early efforts demonstrated the feasibility of decoding visual information from brain
712 activity, laying the groundwork for more advanced approaches. Subsequent progress was marked by
713 the integration of deep learning, particularly through representation alignment techniques. Methods
714 such as BrainCLIP (Liu et al., 2023) and a series of related works (Benchetrit et al., 2023; Xia et al.,
715 2024; Takagi & Nishimoto, 2023b; Qian et al., 2023) leveraged contrastive learning frameworks in-
716 spired by CLIP (Radford et al., 2021) to align neural signals with high-level visual representations.
717 These approaches significantly enhanced decoding accuracy, even under constraints of limited train-
718 ing data.719 A paradigm shift occurred with the introduction of generative frameworks. Mind’s Eye (Scotti et al.,
720 2023) pioneered the application of the unCLIP approach from DALLE-2 (Ramesh et al., 2022) to
721 fMRI-based image reconstruction, incorporating a prior diffusion model to refine neural features
722 before generation. This multi-stage framework was subsequently adapted for EEG-based visual de-
723 coding (Li et al., 2024; Zhang et al., 2025; Xiao et al., 2025), yielding visually impressive results.
724 However, these complex pipelines introduced significant challenges: error propagation across multi-
725 ple processing stages, substantial computational requirements, and difficulty maintaining coherence
726 between the original neural signal and the generated output. These limitations highlight the need for
727 more direct, efficient approaches that preserve the relationship between neural activity and visual
728 reconstruction.729
730 A.2 VISUAL GENERATIVE MODELS
731732 Contemporary visual generative models primarily fall into two paradigms: diffusion models and au-
733 toregressive models, each with distinct computational characteristics particularly relevant for neural
734 decoding applications.735 Diffusion models (Ho et al., 2020; Song et al., 2020; Podell et al., 2023; Peebles & Xie, 2023;
736 Esser et al., 2024) generate images by reversing a gradual noise-addition process, iteratively re-
737 fining random noise into coherent visual content through multiple denoising steps. While these
738 models produce high-quality images with exceptional detail, they require substantial computational
739 resources and typically involve 25-50 sequential denoising steps, limiting their applicability in
740 resource-constrained or real-time scenarios.741 Autoregressive models (Esser et al., 2021; Yu et al., 2022; Van Den Oord et al., 2017; Razavi et al.,
742 2019; Lee et al., 2022) offer an alternative approach by discretizing images into token sequences
743 and predicting each token conditionally on previous ones. Traditionally, these models generated
744 images in raster-scan order (pixel-by-pixel or patch-by-patch), which limited their ability to capture
745 global structure and produce high-resolution outputs (Yu et al., 2022). However, recent innovations
746 in hierarchical autoregressive modeling, particularly Visual Autoregressive modeling (VAR) (Tian
747 et al., 2024), have transformed this landscape by introducing a coarse-to-fine generation strategy
748 called “next-scale prediction.” This approach progressively elaborates visual details across multiple
749 resolution scales, requiring significantly fewer sampling steps (typically 8-12) while maintaining
750 global coherence.751 The efficiency advantages of VAR make it particularly well-suited for neural decoding applications,
752 where computational constraints are significant and the hierarchical nature of the generation process
753 aligns conceptually with theories of visual processing in the brain. AVDE builds upon this paradigm,
754 leveraging its computational efficiency and hierarchical structure to establish a more direct mapping
755 between EEG signals and visual representations.

756 **B HYPERPARAMETER SETTINGS**
757758 Table 5: Hyperparameters for the LaBraM-based EEG encoder.
759

760	761	Hyperparameters	762	Values
762	763	Input channels	764	$\{1,8,8\}$
763	764	Output channels	765	$\{8,8,8\}$
764	765	Temporal Encoder	766	Kernel size $\{15,3,3\}$
765	766		767	Stride $\{8,1,1\}$
766	767		768	Padding $\{7,1,1\}$
767	768		769	
768	769	Transformer encoder layers	770	12
769	770	Hidden size	771	200
770	771	MLP size	772	800
771	772	Attention head number	773	10
772	773		774	
773	774	Batch size	775	1024
774	775	Peak learning rate	776	2e-3
775	776	Minimal learning rate	777	1e-5
776	777	Learning rate scheduler	778	Cosine
777	778	Optimizer	779	AdamW
778	779	Adam β	780	(0.9,0.99)
779	780	Weight decay	781	0.05
780	781	Total epochs	782	50
781	782	Warmup epochs	783	5
782	783	Drop path	784	0.1
783	784	Layer-wise learning rate decay	785	0.8
784	785	Contrastive loss factor λ	786	0.8

786 Table 6: Hyperparameters for the visual autoregressive transformer.
787

788	789	Hyperparameters	790	Values
789	790	Number of layers	791	16
790	791	Hidden size	792	1024
791	792	MLP size	793	4096
792	793	Number of heads	794	16
793	794	Number of scales/steps	795	10
794	795	Size of scales	796	(1, 2, 3, 4, 5, 6, 8, 10, 13, 16)
795	796		797	
796	797	Batch size	798	128
797	798	Peak learning rate	799	2e-5
798	799	Minimal learning rate	800	2e-6
799	800	Learning rate scheduler	801	Cosine
800	801	Optimizer	802	AdamW
801	802	Adam β	803	(0.9,0.95)
802	803	Weight decay	804	0.05
803	804	Total epochs	805	30
804	805	Warmup epochs	806	0
805	806	Drop path	807	0.1
806	807	Normalization ϵ	808	1e-6
807	808	Drop rate	809	0
808	809	Attention drop rate	809	0
809	809	Condition drop rate	809	0.1

810 C RESULTS ON EEG-IMAGENET
811812
813 C.1 DATASET DESCRIPTION
814815 To verify the generalizability of our method, we also conducted experiments on the first 8 subjects
816 from the EEG-ImageNet dataset (Zhu et al., 2024). The visual stimuli consisted of 80 object cate-
817 gories drawn from a subset of ImageNet21k. Each category contained 50 manually curated images,
818 yielding 4,000 EEG-image pairs per subject. After randomizing the category order, the 50 images
819 in each category were sequentially presented using the RSVP paradigm, with each image shown for
820 500 ms. EEG signals were recorded from 62 channels, band-pass filtered between 0.5 and 80 Hz,
821 and sampled at 1,000 Hz. Following the protocol of the original study, we used the first 30 images
822 of each category for training and the remaining 20 for testing.
823
824825 C.2 CLASSIFICATION
826827
828
829 Table 7: Overall accuracy of 80-way classification on EEG-ImageNet. Each cell presents the Top-1
830 accuracy on the first line and the Top-5 accuracy on the second line. For each subject, the highest
831 accuracy values are indicated in bold.
832833
834
835
836
837
838
839
840
841
842
843

Method	Sub-01	Sub-02	Sub-03	Sub-04	Sub-05	Sub-06	Sub-07	Sub-08
EEGNetV4 (Lawhern et al., 2018)	0.297 0.612	0.241 0.607	0.206 0.639	0.139 0.402	0.177 0.553	0.165 0.689	0.152 0.672	0.097 0.357
EEGConformer (Song et al., 2022)	0.301 0.629	0.263 0.618	0.194 0.636	0.141 0.421	0.209 0.552	0.173 0.701	0.151 0.667	0.102 0.365
NICE (Song et al., 2023)	0.293 0.621	0.258 0.607	0.205 0.640	0.138 0.419	0.196 0.561	0.184 0.683	0.149 0.671	0.095 0.359
ATM (Li et al., 2024)	0.306 0.633	0.245 0.624	0.213 0.641	0.151 0.428	0.205 0.557	0.187 0.699	0.146 0.676	0.091 0.353
AVDE (Ours)	0.308 0.634	0.270 0.628	0.227 0.643	0.154 0.435	0.218 0.566	0.329 0.703	0.301 0.695	0.144 0.388

844
845
846 C.3 RECONSTRUCTION
847848
849
850 Table 8: Quantitative assessments of EEG-based visual reconstruction quality on EEG-ImageNet.
851852
853
854
855
856
857
858
859
860
861
862
863

Subject	Low-level		High-level				
	PixCorr \uparrow	SSIM \uparrow	AlexNet(2) \uparrow	AlexNet(5) \uparrow	Inception \uparrow	CLIP \uparrow	SwAV \downarrow
Sub-01	0.106	0.283	0.577	0.664	0.556	0.539	0.626
Sub-02	0.096	0.277	0.564	0.643	0.547	0.527	0.639
Sub-03	0.095	0.264	0.541	0.617	0.525	0.501	0.651
Sub-04	0.080	0.249	0.505	0.578	0.480	0.479	0.677
Sub-05	0.086	0.253	0.528	0.592	0.504	0.482	0.664
Sub-06	0.124	0.318	0.616	0.701	0.603	0.568	0.605
Sub-07	0.119	0.308	0.599	0.694	0.581	0.553	0.613
Sub-08	0.077	0.232	0.480	0.553	0.461	0.457	0.690

Table 9: Reconstruction quality on EEG-ImageNet with baselines.

(a) Li et al. (2024)

Subject	Low-level		High-level				
	PixCorr \uparrow	SSIM \uparrow	AlexNet(2) \uparrow	AlexNet(5) \uparrow	Inception \uparrow	CLIP \uparrow	SwAV \downarrow
Sub-01	0.099	0.269	0.559	0.629	0.529	0.511	0.648
Sub-02	0.092	0.264	0.543	0.611	0.532	0.501	0.659
Sub-03	0.091	0.255	0.525	0.603	0.508	0.480	0.675
Sub-04	0.078	0.236	0.491	0.563	0.466	0.460	0.698
Sub-05	0.081	0.240	0.504	0.569	0.485	0.460	0.692
Sub-06	0.119	0.307	0.589	0.678	0.584	0.545	0.628
Sub-07	0.114	0.296	0.584	0.669	0.565	0.538	0.642
Sub-08	0.074	0.224	0.462	0.529	0.448	0.435	0.699

(b) CognitionCapturer (Zhang et al., 2025)

Subject	Low-level		High-level				
	PixCorr \uparrow	SSIM \uparrow	AlexNet(2) \uparrow	AlexNet(5) \uparrow	Inception \uparrow	CLIP \uparrow	SwAV \downarrow
Sub-01	0.105	0.267	0.551	0.631	0.536	0.509	0.662
Sub-02	0.098	0.263	0.540	0.606	0.539	0.499	0.646
Sub-03	0.088	0.262	0.515	0.593	0.509	0.481	0.679
Sub-04	0.091	0.243	0.494	0.570	0.461	0.456	0.690
Sub-05	0.075	0.244	0.492	0.567	0.483	0.460	0.690
Sub-06	0.113	0.305	0.601	0.691	0.584	0.549	0.616
Sub-07	0.109	0.293	0.595	0.662	0.564	0.553	0.643
Sub-08	0.077	0.213	0.472	0.532	0.449	0.430	0.693

918 **D RESULTS ON THINGS-MEG**
919920 **D.1 DATASET DESCRIPTION**
921923 We also verified our method on the THINGS-MEG dataset (Hebart et al., 2023) to examine its gen-
924 eralizability across neural modalities. The training set consists of 1854 categories, with 12 images
925 in each category, and the test set contains 200 categories. Each image in the dataset was displayed
926 for 500 ms. The preprocessing pipeline includes bandpass filtering of [0.1, 40] Hz, downsampling
927 to 200 Hz and baseline correction. Similar to THINGS-EEG, we ran the retrieval and reconstruction
928 experiments on this dataset.
929930 **D.2 RETRIEVAL**
931934 Table 10: Overall accuracy of 200-way zero-shot retrieval on THINGS-MEG. Each cell presents
935 Top-1 accuracy on the left and Top-5 accuracy on the right. For each subject, the highest accuracy
936 values are indicated in bold.
937

Method	Sub-01	Sub-02	Sub-03	Sub-04
EEGNetV4 (Lawhern et al., 2018)	0.118 / 0.327	0.136 / 0.351	0.171 / 0.382	0.192 / 0.459
EEGConformer (Song et al., 2022)	0.071 / 0.218	0.086 / 0.234	0.119 / 0.259	0.129 / 0.324
NICE (Song et al., 2023)	0.173 / 0.427	0.162 / 0.322	0.185 / 0.487	0.198 / 0.449
ATM (Li et al., 2024)	0.196 / 0.456	0.153 / 0.381	0.235 / 0.512	0.242 / 0.489
AVDE (Ours)	0.221 / 0.498	0.211 / 0.466	0.244 / 0.533	0.266 / 0.494

948 **D.3 RECONSTRUCTION**
949951 Table 11: Reconstruction quality on THINGS-MEG with baselines.
952953 (a) Li et al. (2024)
954

Subject	Low-level		High-level				
	PixCorr \uparrow	SSIM \uparrow	AlexNet(2) \uparrow	AlexNet(5) \uparrow	Inception \uparrow	CLIP \uparrow	SwAV \downarrow
Sub-01	0.119	0.327	0.691	0.802	0.681	0.699	0.642
Sub-02	0.094	0.276	0.661	0.744	0.627	0.673	0.664
Sub-03	0.128	0.315	0.714	0.826	0.716	0.728	0.645
Sub-04	0.066	0.281	0.702	0.754	0.664	0.661	0.677

963 (b) AVDE (Ours)
964

Subject	Low-level		High-level				
	PixCorr \uparrow	SSIM \uparrow	AlexNet(2) \uparrow	AlexNet(5) \uparrow	Inception \uparrow	CLIP \uparrow	SwAV \downarrow
Sub-01	0.121	0.329	0.721	0.804	0.692	0.699	0.658
Sub-02	0.109	0.351	0.732	0.789	0.681	0.712	0.649
Sub-03	0.152	0.344	0.740	0.801	0.668	0.725	0.654
Sub-04	0.113	0.333	0.719	0.808	0.699	0.742	0.641

972 E ADDITIONAL RESULTS ON THINGS-EEG
973
974
975
976977 E.1 RETRIEVAL
978
979980 Table 12: Standard deviation values of Top-1 and Top-5 accuracy on the 200-way zero-shot retrieval
981 task. These results are complement to Table 1.

Method	Sub-01	Sub-02	Sub-03	Sub-04	Sub-05	Sub-06	Sub-07	Sub-08	Sub-09	Sub-10
Within-subject: train and test on one subject										
EEGNetV4 Lawhern et al. (2018)	0.012 0.045	0.008 0.041	0.008 0.025	0.019 0.031	0.020 0.010	0.012 0.031	0.015 0.035	0.015 0.034	0.017 0.026	0.009 0.028
EEGConformer Song et al. (2022)	0.014 0.009	0.015 0.008	0.013 0.044	0.020 0.043	0.017 0.017	0.010 0.015	0.009 0.026	0.011 0.040	0.018 0.018	0.019 0.041
NICE Song et al. (2023)	0.010 0.030	0.014 0.028	0.012 0.037	0.016 0.039	0.012 0.027	0.008 0.043	0.015 0.037	0.031 0.033	0.010 0.044	0.009 0.042
ATM Li et al. (2024)	0.016 0.017	0.006 0.022	0.022 0.015	0.015 0.018	0.011 0.011	0.027 0.024	0.013 0.026	0.022 0.026	0.016 0.023	0.020 0.028
AVDE (Ours)	0.018 0.020	0.011 0.023	0.015 0.016	0.012 0.021	0.017 0.019	0.027 0.025	0.018 0.018	0.023 0.028	0.012 0.019	0.017 0.025
Cross-subject: leave one subject out for test										
EEGNetV4 Lawhern et al. (2018)	0.012 0.025	0.010 0.022	0.009 0.021	0.014 0.023	0.013 0.020	0.011 0.022	0.012 0.024	0.013 0.025	0.011 0.021	0.012 0.022
EEGConformer Song et al. (2022)	0.010 0.009	0.009 0.008	0.009 0.021	0.013 0.022	0.012 0.014	0.009 0.013	0.008 0.017	0.008 0.021	0.011 0.015	0.011 0.020
NICE Song et al. (2023)	0.009 0.019	0.010 0.018	0.009 0.024	0.011 0.025	0.010 0.017	0.008 0.021	0.011 0.023	0.017 0.020	0.008 0.025	0.007 0.023
ATM Li et al. (2024)	0.013 0.015	0.008 0.016	0.015 0.013	0.011 0.015	0.010 0.011	0.018 0.017	0.011 0.018	0.015 0.018	0.013 0.016	0.014 0.019
AVDE (Ours)	0.014 0.016	0.010 0.017	0.012 0.013	0.011 0.015	0.013 0.014	0.019 0.018	0.014 0.015	0.017 0.020	0.011 0.015	0.013 0.018

1003
1004 E.2 RECONSTRUCTION
1005
10061007 Table 13: Quantitative assessments of EEG-based visual reconstruction quality on all subjects.
1008
1009

Subject	Low-level			High-level			
	PixCorr \uparrow	SSIM \uparrow	AlexNet(2) \uparrow	AlexNet(5) \uparrow	Inception \uparrow	CLIP \uparrow	SwAV \downarrow
Sub-01	0.138	0.362	0.750	0.837	0.725	0.721	0.609
Sub-02	0.120	0.388	0.754	0.811	0.710	0.737	0.604
Sub-03	0.168	0.379	0.763	0.827	0.704	0.746	0.603
Sub-04	0.126	0.361	0.748	0.835	0.731	0.765	0.593
Sub-05	0.145	0.373	0.743	0.814	0.691	0.732	0.586
Sub-06	0.128	0.364	0.768	0.832	0.707	0.717	0.584
Sub-07	0.126	0.351	0.760	0.827	0.733	0.731	0.588
Sub-08	0.188	0.396	0.817	0.889	0.765	0.795	0.557
Sub-09	0.156	0.339	0.753	0.810	0.713	0.726	0.600
Sub-10	0.173	0.349	0.807	0.871	0.764	0.800	0.543

Table 14: Reconstruction quality on THINGS-EEG with baselines.

(a) Li et al. (2024)

Subject	Low-level		High-level				
	PixCorr \uparrow	SSIM \uparrow	AlexNet(2) \uparrow	AlexNet(5) \uparrow	Inception \uparrow	CLIP \uparrow	SwAV \downarrow
Sub-01	0.136	0.354	0.720	0.831	0.710	0.728	0.589
Sub-02	0.108	0.301	0.688	0.775	0.657	0.705	0.618
Sub-03	0.145	0.340	0.747	0.858	0.751	0.759	0.590
Sub-04	0.075	0.305	0.729	0.782	0.694	0.688	0.623
Sub-05	0.133	0.352	0.738	0.791	0.637	0.707	0.612
Sub-06	0.140	0.363	0.753	0.834	0.686	0.742	0.595
Sub-07	0.155	0.360	0.769	0.851	0.684	0.697	0.608
Sub-08	0.163	0.345	0.786	0.868	0.730	0.770	0.575
Sub-09	0.134	0.338	0.736	0.788	0.626	0.652	0.606
Sub-10	0.112	0.329	0.735	0.843	0.674	0.708	0.570

(b) CognitionCapturer (Zhang et al., 2025)

Subject	Low-level		High-level				
	PixCorr \uparrow	SSIM \uparrow	AlexNet(2) \uparrow	AlexNet(5) \uparrow	Inception \uparrow	CLIP \uparrow	SwAV \downarrow
Sub-01	0.148	0.334	0.741	0.626	0.666	0.711	0.592
Sub-02	0.147	0.344	0.764	0.618	0.661	0.725	0.590
Sub-03	0.140	0.307	0.715	0.549	0.690	0.710	0.603
Sub-04	0.126	0.355	0.801	0.660	0.701	0.765	0.543
Sub-05	0.130	0.343	0.731	0.639	0.594	0.655	0.611
Sub-06	0.122	0.337	0.748	0.620	0.646	0.688	0.630
Sub-07	0.145	0.355	0.777	0.623	0.731	0.721	0.576
Sub-08	0.175	0.366	0.760	0.610	0.721	0.744	0.577
Sub-09	0.148	0.337	0.731	0.623	0.625	0.692	0.605
Sub-10	0.152	0.389	0.773	0.664	0.657	0.736	0.569

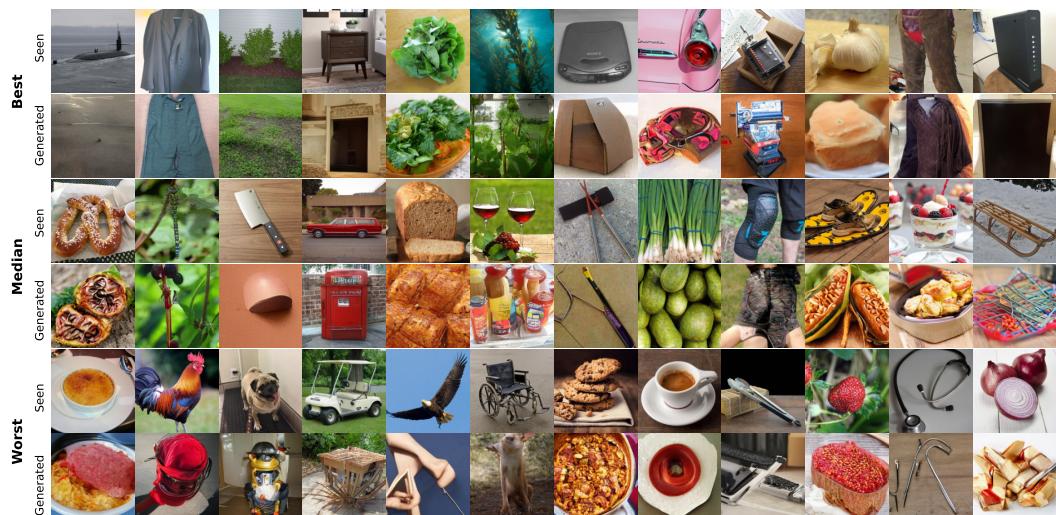


Figure 6: Subject-08's Best, Medium, and Worst reconstructions selected based on PixCorr.



Figure 7: Subject-08’s Best, Medium, and Worst reconstructions selected based on SwAV.

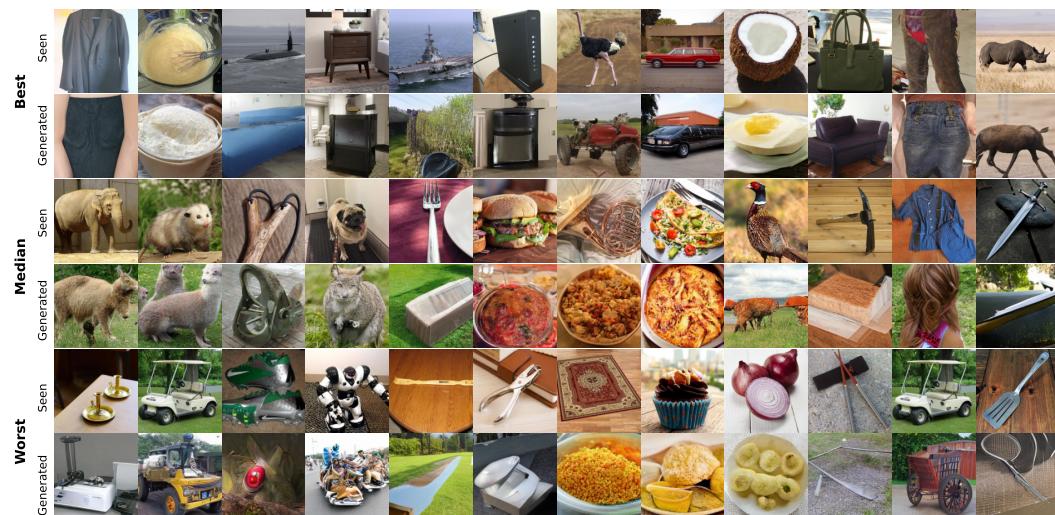


Figure 8: Subject-10's Best, Medium, and Worst reconstructions selected based on PixCorr.



Figure 9: Subject-10's Best, Medium, and Worst reconstructions selected based on SwAV.

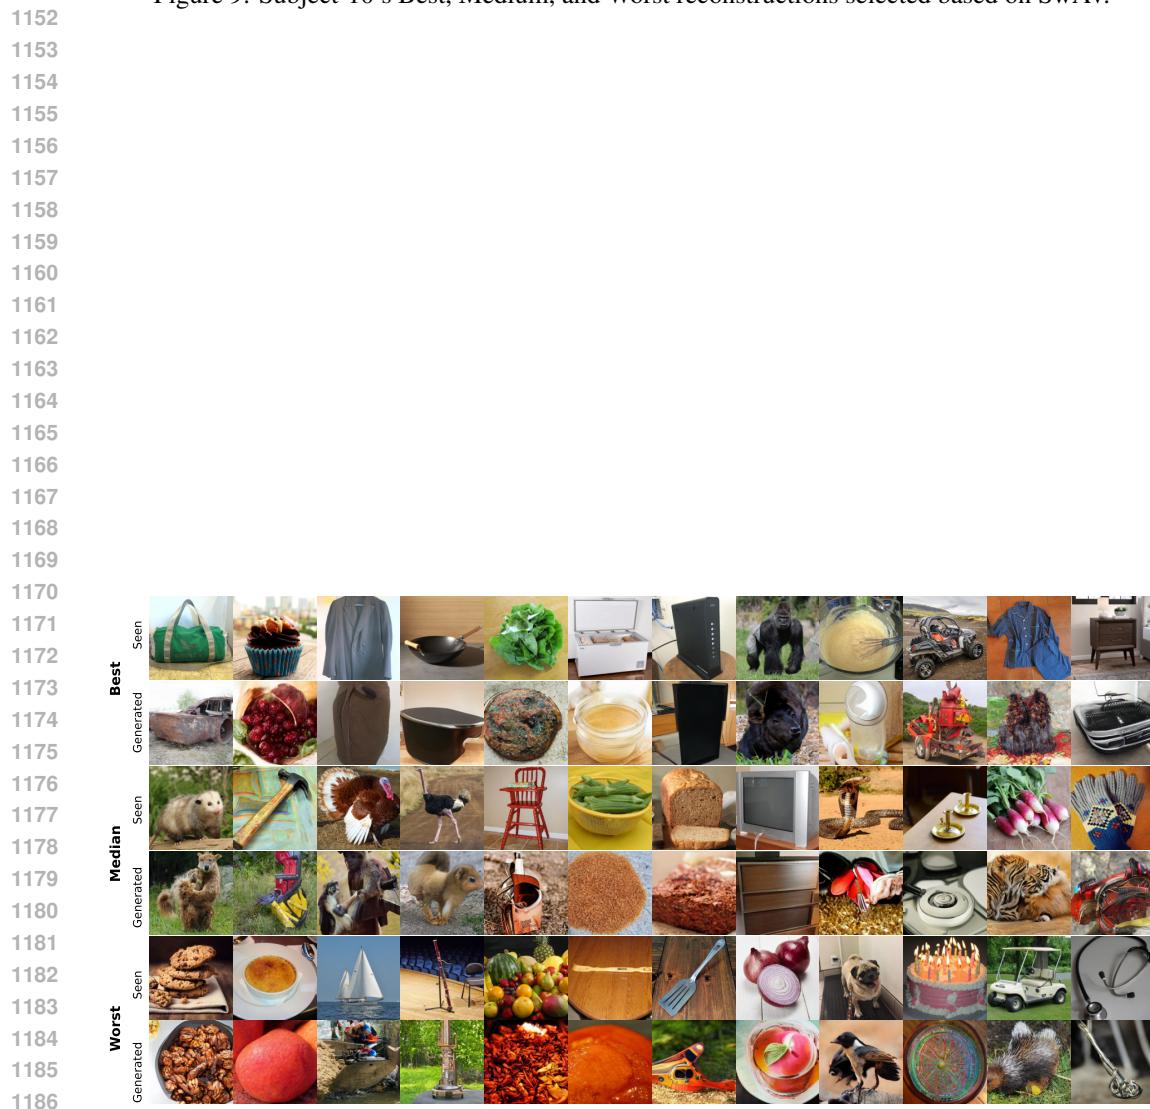


Figure 10: Subject-03's Best, Medium, and Worst reconstructions selected based on PixCorr.



Figure 11: Subject-03’s Best, Medium, and Worst reconstructions selected based on SwAV.