

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 LEARNING FREQUENCY DOMAIN CODES FOR SEMANTIC VISION

Anonymous authors

Paper under double-blind review

## ABSTRACT

Visually semantic concepts such as objects and categories provide a natural foundation for semantic reasoning, yet standard deep learning-based vision models routinely extract and aggregate features using homogeneous stacks of spatial layers. As a result, feature representations are learnt implicitly without clear organisation, rendering decision-making processes opaque and difficult to interpret. Psychovisual processing provides a way to mimic how the brain encodes and interprets visual information that produces higher abstractions from low-level processing. In this paper, we propose Semantic Visual Coding (SVC), a learnt frequency domain representation that introduces explicit psychovisual abstraction into convolutional neural networks (CNNs). Inspired by psychovisually motivated image codes from the 1990s, SVC learns band-limited filters that encode task-relevant semantics as distinct regions of the frequency domain. These converge towards sparse (data-driven) coronal patterns that suggest a natural representation scheme for semantic abstractions supporting model reasoning. We also introduce a framework that adapts CNNs to be psychovisually aware by combining traditional low-level spatial feature extraction with high-level abstraction in the frequency domain via SVC, which we call ‘PsychoNet’. Salience analyses show that PsychoNet’s spatial layers extract highly interpretable object parts and morphological features, unlike blob-like regions produced by standard CNNs. It further finds that SVC forms structured selections of these parts that are organised by spatial scale, suggesting frequency domain abstraction as a promising direction for interpretable models which reveal the semantic features they employ.

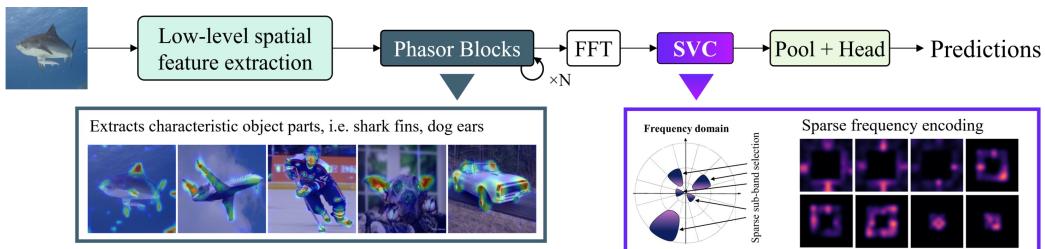


Figure 1: The brain encodes and interprets visual information using **psychovisual processing**, which separates feature extraction from higher cognition using intermediate abstractions. **PsychoNet** introduces similar pipelines to CNNs. Early spatial layers extract low-level features, similar to early cortical processing, and **Phasor Blocks** localise key characteristic object parts. Subsequently, these parts are encoded in the frequency domain by **Semantic Visual Coding (SVC)** into sparse frequency sub-bands. We believe that this is a naturally emergent representation for semantic information, similar to psychovisual abstractions. FFT denotes the Fast Fourier Transform.

## 1 INTRODUCTION

Ever since the ImageNet challenge popularised deep learning for computer vision (Krizhevsky et al., 2017; Deng et al., 2009), architectural advances have focused on the design of spatial domain feature extractors, from convolution layers (He et al., 2016; Xie et al., 2017; Huang et al., 2017; Liu

et al., 2022) to more recent token mixers based on attention mechanisms (Dosovitskiy et al., 2021; Tolstikhin et al., 2021; Gao et al., 2021; Rao et al., 2023). Although these models achieve impressive performance, the way they reason in deeper layers is often opaque and difficult to interpret. Psychovisual processing—the way human vision encodes and interprets visual information—separates feature extraction from higher cognition using intermediate abstractions, like objects, relations, and categories, providing a natural basis for reasoning (Quiroga et al., 2005; Kriegeskorte et al., 2008; Quiroga, 2012; Le et al., 2024). In contrast, current architectures refine and aggregate image features using homogeneous stacks of spatial layers, with the feature processing emerging implicitly through training and often lacking clear organisation.

In this work, we propose Semantic Visual Coding (SVC), a frequency-domain high-level processing module that bridges low-level feature extraction and decision making in CNNs, revealing interpretable, object-part-based semantic representations that underpin model reasoning within a psychovisual-inspired pipeline. Concretely, they are implemented with data-driven band-limited frequency filters, inspired by psychovisual coding schemes from the 1990s that targeted perceptually salient frequencies found by human vision studies (Saadane et al., 1994; Guedon et al., 1995; Saadane et al., 1998). Semantic Visual Coding (SVC) extends this idea to high-level abstraction by allowing networks to discover by learning the sparse frequency subsets most relevant to a given task. These promote richer intermediate abstractions grounded in task-oriented semantic features and yield potentially interpretable views of the model’s reasoning, pointing to a naturally emergent framework for high-level abstraction with semantic reasoning. In future, we also aim to apply SVC to domains where the frequency domain is the natural measurement space, particularly magnetic resonance imaging (MRI) (Chandra et al., 2021) as transparency and trustworthy model behaviour are crucial due to the high-stakes nature of medical applications.

Additionally, we develop PsychoNet, a framework that incorporates SVC into conventional convolutional neural networks (CNNs), demonstrated on the widely used ResNet and state-of-the-art ConvNeXt architectures. PsychoNet establishes a coherent dual-domain pipeline (Figure 1): initial low-level image features are extracted in the spatial domain and augmented by learning complex-valued representations, then SVC constructs high-level abstractions in the frequency domain that support decision making with semantic reasoning. To the best of our knowledge, this work presents the first data-driven exploration of the frequency domain for high-level representation learning in vision, whereas prior studies focus mainly on lower level feature learning or parameterising spatial models (Chi et al., 2020; Rippel et al., 2015; Rao et al., 2023). To summarise, the key contributions of this work are:

- Inspired by psychovisual abstraction, we introduce SVC, a deep-learning based module that automatically learns frequency domain representations of high-level semantic image information for a given vision task. These emerge as sparse selections of coronal frequency sub-bands in the discrete Fourier Transform (DFT).
- Our PsychoNet (Figure 1) is a framework that integrates SVC into conventional CNN models. We demonstrate it enables interpretable psychovisual-like processing on ResNet and ConvNeXt architectures while maintaining or improving performance across various classification tasks.
- Through salience analysis, we demonstrate clear evidence of semantic reasoning by revealing that intermediate spatial layers consistently focus on meaningful object parts, which SVC encodes using data-driven, psychovisual coding-inspired filters. Abstractions produced by SVC are shown to form structured encodings of the object parts, whose use in final decision making mirrors human psychovisual processing and highlights a promising pathway toward interpretable model reasoning.

## 2 BACKGROUND

In this section, we review related prior computer vision works on methods with biological motivations, as well as those that use the frequency domain. Additional details/background about psychovisual coding, the Fourier Transform and complex-valued neural networks are provided in Appendix A.

108 **Biologically Inspired Vision.** Biologically inspired approaches in computer vision predominantly  
 109 focus on modelling early vision stages. In particular, much attention has been given to receptive  
 110 fields (RF) - regions of visual stimuli that elicit strong neural responses in the visual cortex. Mammalian  
 111 RFs are known to act as directional differential operators, closely resembling traditional im-  
 112 age processing functions like wavelets and Gabor filters (Olshausen & Field, 1996; Hubel & Wiesel,  
 113 1962; Ringach, 2002). These parallels motivated their use in approximating low-level human vi-  
 114 sion, serving as effective feature extractors for basic visual structures like edges and shapes. In deep  
 115 learning, these functions have been used to build neural networks that mimic cortical pathways (Liu  
 116 et al., 2023), and early-layer CNN kernels also perform similar directional operations (Krizhevsky  
 117 et al., 2017; Rippel et al., 2015). Beyond RFs, cortical responses have also been modelled from a  
 118 frequency domain perspective.

119 Research conducted in the 1990s by French researchers led by Dominique Barba in understanding  
 120 human vision included psychovisual experiments determining frequency sensitivities of the human  
 121 visual cortex (Saadane et al., 1994; Senane et al., 1995; Saadane et al., 1998). These informed the  
 122 design of psychovisual coding, an image quantization and compression scheme that is perceptually  
 123 lossless to humans. It first decomposes the frequency domain into a number of coronal sub-bands  
 124 (Figure 2 (left)), then for each sub-band applies specific quantization thresholds derived based on the  
 125 discovered sensitivities. This enabled only frequencies corresponding to perceptually salient image  
 126 figures to be encoded in full, while the rest are removed or heavily compressed without affecting  
 127 perceived image quality. An extended review covering the motivations and additional background  
 128 for psychovisual coding is included in Appendix A.1.

129 **Frequency Domain Learning.** Frequency analysis has long been a staple in traditional image pro-  
 130 cessing. Unlike the spatial domain, which is highly localised and expresses features in contiguous  
 131 pixel neighbourhoods, the frequency domain is more conducive to global representations (see Ap-  
 132 pendix A.2). Formulated in this space, image processing functions like ridgelets (Candés & Donoho,  
 133 1999), curvelets (Starck et al., 2002) and contourlets (Do & Vetterli, 2005) have appealing sparse  
 134 representations. In fact, they bear a strong resemblance to psychovisual codes since they target  
 135 specific selections of sub-bands, corresponding to features from different spatial scales. Although  
 136 these functions have been incorporated into neural networks before, they are only effective on small  
 137 problems due to their handcrafted nature (Liu et al., 2021).

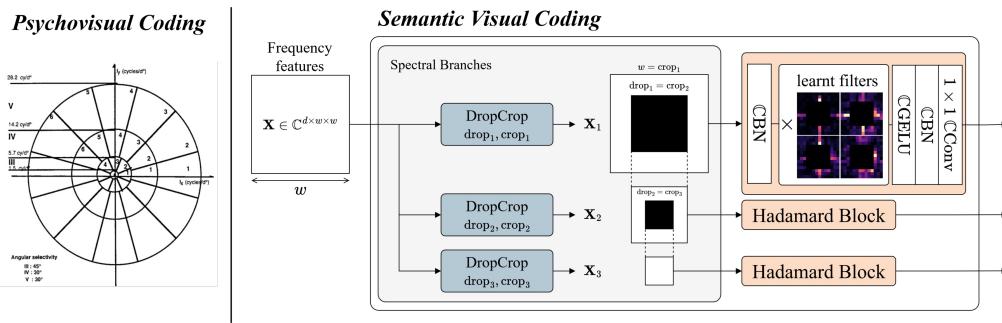
138 In deep learning, the frequency domain has primarily been used to exploit the Convolution Theorem  
 139 (Gonzalez & Woods, 2014), whereby spatial circular convolution becomes simple elementwise mul-  
 140 tiplication in the frequency domain. Many works leveraging this property are performance-driven:  
 141 (Li et al., 2020a; Chi et al., 2020; Guan et al., 2021) use frequency-domain filters to accelerate  
 142 CNNs and incorporate global context, while (Rao et al., 2023; Lee-Thorp et al., 2021; Huang et al.,  
 143 2023) employ global frequency filters as lightweight and effective token mixers for transformer-  
 144 style models. Other studies use frequency-domain re-parameterizations of CNNs to analyse model  
 145 properties and behaviours (Rippel et al., 2015; Grabinski et al., 2023; Kabri et al., 2023), such as  
 146 optimal kernel structures. In contrast, our work contributes to frequency-domain representation  
 147 learning, a direction that remains comparatively underexplored. Since global context is crucial for  
 148 high-level features (Rao et al., 2023; Dosovitskiy et al., 2021), the frequency domain provides a nat-  
 149 ural setting in which to represent and process semantic structure. Our SVC module employs learn-  
 150 able band-limited frequency filters—data-driven analogues of hand-crafted visual codes—to encode  
 151 task-relevant semantic information. Unlike prior works that integrate frequency filters directly into  
 152 their base computational blocks, SVC acts as an abstraction layer bridging feature extraction and  
 153 decision. Another related work uses frequency filters to select and amplify domain-transferable fre-  
 154 quency (Lin et al., 2023), similar to how SVC selects task-relevant features, but their approach is  
 155 not framed as representation learning nor is it used for interpretability. Moreover, SVC also enables  
 the novel interpretable psychovisual-like processing achieved by PsychoNet.

### 3 METHOD

160 **Semantic Visual Coding.** A  $N \times N$  digital image, or a spatial feature map derived from it by  
 161 a neural network, can be viewed as a 2D discrete signal  $x[m, n]$ ,  $m, n \in 0, \dots, N$ . This can be  
 represented in the frequency domain as a linear combination of complex-valued sinusoids via the

162 2D DFT(Cooley et al., 1969):  
163

164 
$$X[u, v] = \frac{1}{N^2} \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} x[m, n] e^{-2\pi i (\frac{um+vn}{N})} \quad (1)$$
  
165

166 where  $i$  denotes the imaginary unit. These weights are the (frequency) spectrum of the image and is  
167 a complex-valued space known as the frequency domain, which can be computed efficiently using  
168 the Fast Fourier Transform (FFT) (Cooley et al., 1969). Psychovisual coding (Saadane et al., 1998)  
169 partitions this space into radial sub-bands (2 (left)), and assigns each a threshold corresponding  
170 to sensitivity to human vision. These thresholds decide the level of granularity when quantizing  
171 images, so that perceptually important features are preserved while others are coarsely represented  
172 or discarded.184 Figure 2: **(Left)** Hand-crafted psychovisual coding from Saadane et al. (1998), which quantizes  
185 perceptually salient radial frequencies determined by human vision experiments. **(Right)** Our Se-  
186 mantic Visual Coding module, a data-driven adaptation of psychovisual coding. It uses (1) Spectral  
187 Branches for radial spectral decomposition (2) Hadamard Blocks to apply learnt element-wise filters  
188 and channel mixing. CCConv/BN/GELU denote complex-valued convolution, batch norm and GELU  
189 operations - see Appendix B.2190 We introduce *Semantic Visual Coding* (Figure 2 (right)) which aims to generalize this coding principle  
191 beyond low-level vision and adapt it to high-level features in deep network layers. In this setting,  
192 the selection of frequencies should no longer be fixed by handcrafted thresholds, but instead learnt  
193 directly from data to encode task-relevant semantic information. Semantic Visual Coding has the  
194 following formulation:195 Let  $\mathbf{X} \in \mathbb{C}^{d \times w \times w}$  be frequency domain input features, where  $d$  is the number of channels and  
196  $w \times w$  the spatial size.

1. We apply *Spectral Branches* which replicate the radial frequency partitioning in psychovisual codes. These divide  $\mathbf{X}$  into disjoint rectangular sub-bands  $\mathbf{X}_1, \mathbf{X}_2, \dots$  using *DropCrop* blocks, which set a lower frequency boundary ( $\text{drop}_i$ ) by zeroing central frequencies and an upper boundary ( $\text{crop}_i$ ) by cropping  $\mathbf{X}$  to size  $d \times \text{crop}_i \times \text{crop}_i$ .
2. For each sub-band  $\mathbf{X}_i$ , *Hadamard Blocks* apply a set of learnt filters  $\mathbf{W}_i \in \mathbb{C}^{d \times \text{crop}_i \times \text{crop}_i}$  via element-wise (Hadamard) multiplication. Additionally, we also apply Softmax across the channels of  $\mathbf{W}_i$  to amplify important frequency selections and suppress unimportant ones, emulating the quantization in psychovisual coding.
3. Hadamard Blocks further apply complex  $1 \times 1$  convolution block to mix together different information extracted by each channel/filter, yielding our final representations.

209 In practice, for all models we apply Spectral Branches at a spatial resolution of  $w = 14$  and use  
210 three sub-bands with  $[\text{crop}_i, \text{drop}_i]$  values of  $[14, 8], [8, 4]$  and  $[4, 1]$  respectively. More details can  
211 be found in Appendix B.212 **PsychoNet.** The PsychoNet framework (Figure 3) adapts standard spatial CNNs to use Semantic  
213 Visual Coding. This setting enables experimentation to assess if our codes produce meaningful ab-  
214 stract representations that support interpretable semantic reasoning, as well as practical performance  
215 evaluation against standard baselines.

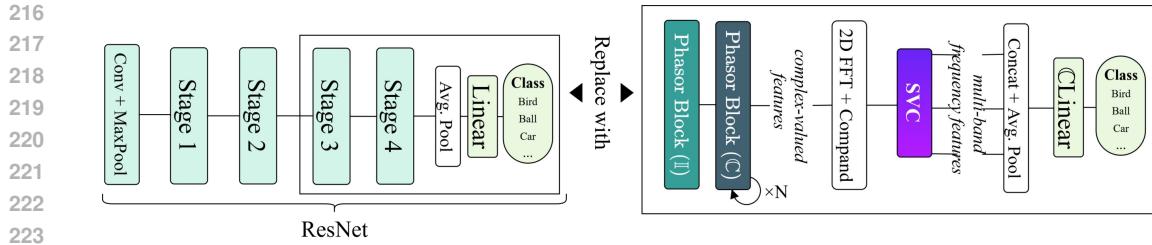


Figure 3: Example: Converting ResNet50/101 image classification models with PsychoNet.

In our experiments, we apply PsychoNet to ResNet (He et al., 2016) and ConvNeXt (Liu et al., 2022) architectures. Below we summarise the main steps for implementing PsychoNet below while full architectural configurations are provided in Appendix B.

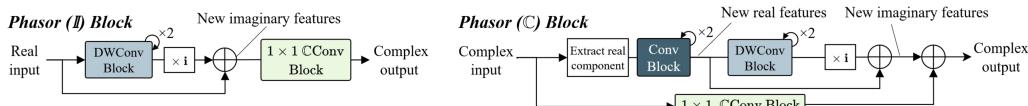
1. A number of low-level feature extraction layers are retained from the base CNN, for example, the first two resolution stages in the case of ResNet50 and ResNet101.
2. The remaining spatial layers are replaced with Phasor Blocks, described further below. Compared to the original CNNs layers, Phasor Blocks typically use higher spatial resolution and only downsample down to  $14 \times 14$  instead of  $7 \times 7$ . Though this increases FLOPs (Appendix C), we found there is insufficient granularity at  $7 \times 7$  to clearly separate low and high frequencies after FFT.
3. 2D FFT is applied to convert features from the spatial to frequency domain. As with most visual features, the magnitude of  $X$ 's DC (0 frequency) and low frequency features typically dominate over those of high frequency ones, so we use a simple companding operation to reduce this imbalance (Appendix B).
4. SVC is applied, and its outputs from each frequency band aggregated. These are then used for output prediction directly in the frequency domain using a complex-valued linear layer.

The ConvNeXt-based PsychoNet differs from its ResNet-based counterpart in only two key architectural aspects: (1) the use of ConvNeXt's  $4 \times 4$  patch embedding layer at the input, and (2)  $7 \times 7$  depthwise convolution blocks are used instead of regular convolution layers within Phasor Blocks, matching ConvNeXt's primary computational block. Despite their minimal nature, these adaptations are sufficient to recover performance close to ConvNeXt, as demonstrated in Section 4.

The following section outlines the motivation and formulation of Phasor Blocks, with additional architectural details provided in Appendix B

**Phasor Blocks.** Filtering in the frequency domain is powerful as it captures information encoded as both magnitude and phase. However, PsychoNet operates on real-valued inputs (natural images), and real features incur the conjugate symmetry of the Fourier Transform (FT), rendering half of the frequency domain redundant. While this constraint matters little when the frequency domain is used solely for convolution (Rao et al., 2023; Li et al., 2020a), it limits learnt filtering from fully exploiting complex representations. To address this, we introduce Phasor Blocks (Figure 4) to augment real-valued spatial features with complementary complex-valued ones, breaking conjugate symmetry and improving the specificity of learned sub-bands. In practice, imaginary components are generated from existing real features using lightweight depthwise convolution based blocks which decouple spatial and channel mixing to encourage cross-channel interaction without altering spatial structure. In natural complex signals, real and imaginary components convey complementary information at the same spatial location (Gonzalez & Woods, 2014; Lee et al., 2022), making it important that the generated imaginary features do not introduce substantial new spatial information. As demonstrated in later sections, this design enables Phasor Blocks to extract meaningful object parts that form the basis of SVC abstractions supporting interpretable semantic reasoning. Further evidence from the ablations in Appendix D.1 and the salience maps in Figure A.10 shows that while SVC alone encourages attention to object parts, the inclusion of Phasor Blocks and their complex representations leads to substantially clearer and more localised part-based salience, enabling greater interpretability.

270  
271  
272  
273  
274  
275  
276  
277  
Figure 4 provides an overview of the two Phasor Block configurations used in this work, with full  
architectural diagrams and design details presented in Appendix B.3.



278  
279  
280  
281  
282  
Figure 4: **Phasor Blocks architectures.** Phasor (I) blocks generate complementary imaginary  
features for real-valued input. Phasor (C) blocks generate additional complex features for complex-  
valued input based on its real component. Implementations of normal convolution (Conv), depthwise  
convolution (DWConv) (Liu et al., 2022) and complex-valued convolution (CConv) blocks for each  
of our models are presented in Figure A.2

## 283 284 4 EXPERIMENTS AND RESULTS 285

286  
287 We applied PsychoNet across multiple ResNet architectures, which provide straightforward and well  
288 understood baselines, as well as ConvNeXt-S, as an example of a state-of-the-art CNN. Quantitative  
289 evaluation was conducted on image classification across small to large-scale datasets: CIFAR-  
290 10/100 (Krizhevsky, 2009) both contain  $\sim$ 50K low-resolutions images, while ImageNet-100 is a  
291 moderate sized subset ( $\sim$ 130K images) of ImageNet (Deng et al., 2009). Finally, we also use the  
292 standard large ImageNet-1K subset containing  $\sim$ 1.2 million training and  $\sim$ 50K validation images.  
293 Full dataset, training and hardware details are presented in Appendix C, and full model configura-  
294 tions in Appendix B.

295 Model	296 Param. (M)	297 # Layers	298 CIFAR-10	299 CIFAR-100	300 IN100	301 IN1K
ResNet50	25.56	54	94.14	78.10	80.90	76.04
Psycho-S	25.35	$65 \uparrow 20.4\% \text{ more}$	<b>95.08</b>	<b>78.97</b>	<b>82.50</b>	<b>76.86</b>
ResNet101	44.55	105	93.64	79.13	81.90	78.43
Psycho-B	42.01	$93 \downarrow 11.4\% \text{ less}$	<b>94.99</b>	<b>79.49</b>	<b>83.60</b>	<b>78.85</b>
ResNet152	60.10	156	93.17	77.51	83.60	79.59
Psycho-L	61.28	$93 \downarrow 40.4\% \text{ less}$	<b>94.95</b>	<b>79.64</b>	<b>84.82</b>	<b>79.85</b>
ResNet270	89.60	276	76.51	50.87	83.80	80.01
Psycho-H	88.61	$93 \downarrow 66.3\% \text{ less}$	<b>94.68</b>	<b>79.89</b>	<b>85.00</b>	<b>80.45</b>
ConvNeXt-S	50.22	113	94.09	76.96	<b>86.98</b>	<b>80.78</b>
PsychoDW	49.51	$106 \downarrow 6.2\% \text{ less}$	<b>95.46</b>	<b>79.67</b>	86.76	80.59

302  
303  
304  
305  
306  
307  
Table 1: Summary of classification results (% top-1 accuracies) on CIFAR-10, CIFAR-100,  
308 ImageNet-100 (IN100) and ImageNet-1K (IN1K). Each pair of rows (separated by horizontal lines)  
309 compares a baseline CNN and the PsychoNet based on it. Further detailed results are presented in  
310 Appendix C.  
311

312  
313 A common characteristic of newer CNNs like ConvNeXt is their reduced dependence on depth for  
314 model scaling (Liu et al., 2022; Xie et al., 2017), whereas large ResNet architectures heavily add  
315 additional high-level layers to increase representational capacity (He et al., 2016). We hypothe-  
316 size that since high-level processing in PsychoNet is handled by our frequency domain modules,  
317 it should also be much less depth-dependent than ResNet. The Psycho-S/B/L/H models, counter-  
318 parts to ResNet50/101/152/270, were designed to test this and do not increase Phasor Block depth  
319 beyond the Psycho-B/ResNet101 size. Instead, parameter parity is maintained by widening (increas-  
320 ing number of feature channels) the existing Phasor Blocks and SVC filters—an equally simple, if  
321 not simpler, scaling strategy than the depth expansion used in ResNet. Table A.2 compares the  
322 resulting channel-width configurations across these model sizes.

323  
Overall, we found that PsychoNet slightly improves the performance of each baseline ResNet across  
324 all datasets, despite Psycho-L and Psycho-H using  $\sim 1.7 \times$  and  $\sim 3 \times$  less layers than their ResNet

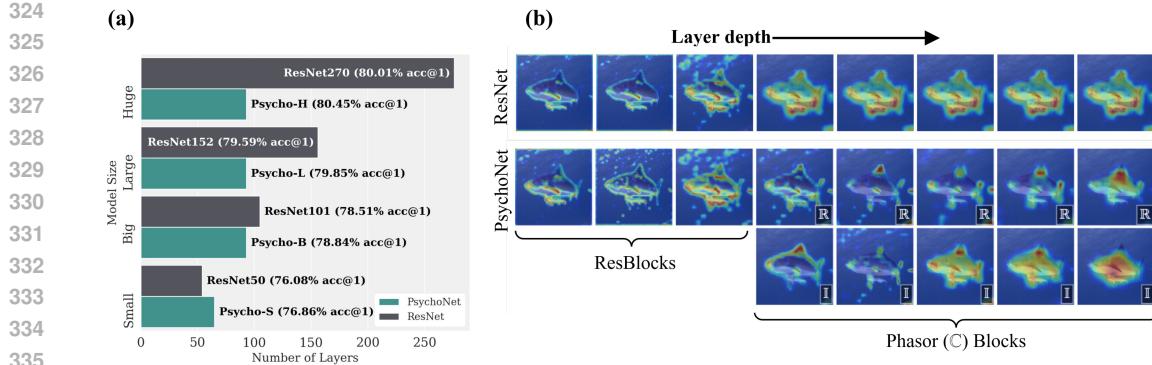


Figure 5: (a) Comparison of model depth when scaling ResNet vs. PsychoNet. (b) Comparison between activation maps (via KPCA-CAM) of Psycho-B and ResNet101 for a range of layer depths. Real and imaginary components are denoted by  $\mathbb{R}$  and  $\mathbb{I}$ .

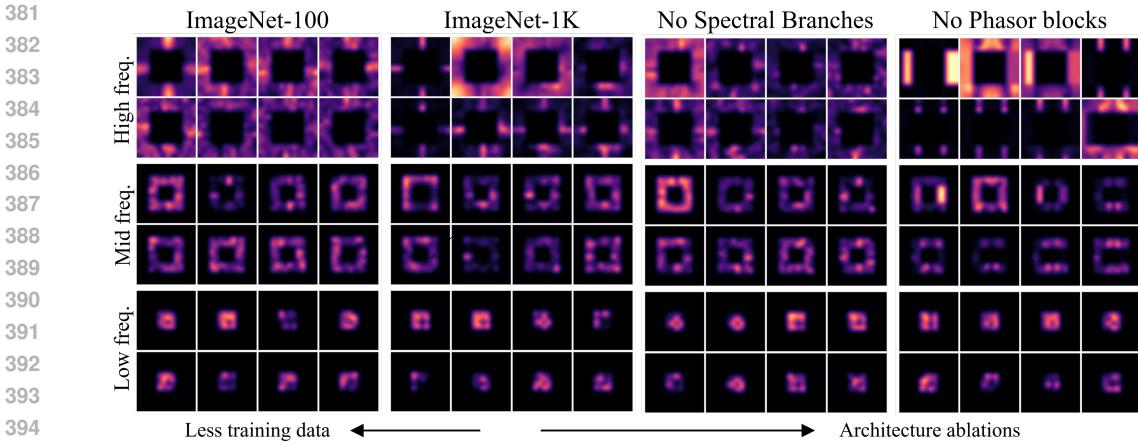
baselines respectively (Table 1, Figure 5 (a)). This demonstrates that SVC is able subsume the role of a significant portion of deep spatial layers, showing it is an effective high-level processing module. We believe this arises because SVC operates on globally focused frequency-domain representations, in contrast to spatial features which are sparse and locally structured, requiring significant depth to progressively aggregate features. We also note that ResNet152 and ResNet270 underperform smaller ResNets on CIFAR-10/100, likely as their large depth is unsuited for the low-resolution images in these datasets. In comparison, our much shallower Psycho-L/H models were unaffected.

On ConvNeXt-S, PsychoDW achieves comparable performance despite its Phasor Block being only a simple adaptation of the ResNet-oriented one, indicating that further targeted adaptations may yield even stronger performance. This demonstrates that PsychoNet extends naturally to modern CNN architectures beyond ResNet. As shown in the following section, PsychoNet also adds interpretable, psychovisual-style reasoning to both ResNet- and ConvNeXt-based models while maintaining performance, which will be particularly beneficial for future applications such as medical image analysis that require transparent and trustworthy models.

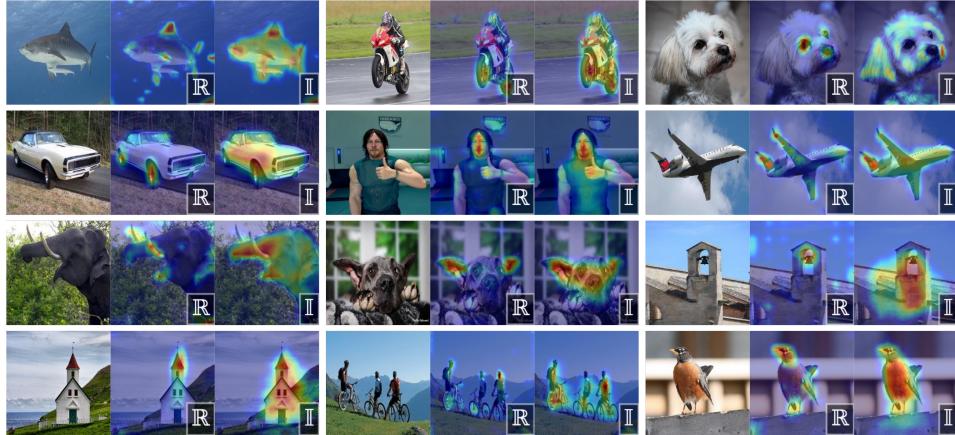
Finally, we also acknowledge that our PsychoNet models use considerably more FLOPs than their respective CNN baselines (Table A.12). The increased computation is attributed to (1) Phasor Blocks requiring higher-resolution features than the CNN layers, to support  $14 \times 14$  SVC filters and (2) complex-valued operations (complex convolution etc.) being poorly optimised in deep learning frameworks. Regarding (1), Appendix D.1.1 presents FLOP-efficient variants of Psycho-S/B that adds an early downsampling step to  $7 \times 7$  and extracts the  $[14, 8]$  sub-bands from intermediate Phasor Block features. On ImageNet-100, this preserves Psycho-S/B performance while matching the FLOPs of ResNet-50/101. While these findings suggest that efficient PsychoNet variants are achievable, our current focus remains on establishing the first instance of interpretable psychovisual separation of low and high-level processing in CNNs, with supporting results presented in the following section.

**Filter learning.** Figure 6 visualises SVC filters learnt by ResNet-based PsychoNet, showing the top spatial principal components as an approximation of the most important frequency features. We find that filters across every sub-band learn very sparse selections of frequencies; similar results were found for the ConvNeXt-S based PsychoDW and are shown in Figure A.5 (a). However, this requires sufficiently large training corpora - the ImageNet-100-trained filters are noticeably noisier than for ImageNet-1K, and those for CIFAR-10/100 (Figure A.4) even more still. This suggests that these patterns correspond to a data-driven representation naturally emergent from visual information. Additionally, we ran ablation experiments to evaluate the effects of Phasor Blocks and Spectral Branching (Appendix D.1). Removing Phasor Blocks (and replacing them with ResNet-style residual bottleneck blocks (ResBlocks)) removes complex-valued spatial features and introduces conjugate symmetry to the frequency domain. This yields symmetric filter features that are far less expressive. Likewise, removing the spectral decomposition of Spectral Branches (we use one global

378 filter instead of three band-limited ones) also reduces filter sparsity. This is likely as exposure to the  
 379 entire frequency domain makes it harder for filters to specialise to specific sub-bands.  
 380



394 Figure 6: SVC filters learnt by Psycho-B trained on ImageNet-100 and ImageNet-1K, as well as for  
 395 two ablation models on ImageNet-1K. Bilinear smoothing has been applied. ‘High/mid/low freq.’  
 396 refer to the [14, 8], [8, 4] and [4, 1] frequency sub-bands created by Spectral Branches. ‘No Spectral  
 397 Branches’ removes Spectral Branches and uses a single Hadamard Block with global filters - we  
 398 extract sub-bands only for the visualisation. ‘No Phasor Blocks’ replaces all Phasor Blocks with  
 399 ResBlocks.  
 400



417 Figure 7: Assorted activation maps (via KPCA-CAM) for mid-level Phasor Blocks of Psycho-B.  
 418 Real and imaginary components are denoted by  $\mathbb{R}$  and  $\mathbb{I}$ . An equivalent visualisation for PsychoDW  
 419 is shown in Figure A.6.  
 420

421 **Representation Analysis.** As our quantitative results suggested SVC likely subsumes the high-  
 422 level processing of deep ResNet layers, we visualise layer activations using KPCA-CAM (Karmani  
 423 et al., 2024) to compare spatial processing between the two models (Figure 5 (b)). This approach  
 424 generates salience maps by projecting activations onto the first principal component of their kernel  
 425 PCA. ResNet’s early layers target low-level features (edges), but later salience regions quickly grow  
 426 to cover the entire subject and it is not particularly clear on which parts of the shark each layer  
 427 is focusing. This likely reflects ResNet’s use of homogenous layer stacks, which, without explicit  
 428 structure, may learn layers that perform diffuse and weakly organised operations. In contrast, early-  
 429 mid level Phasor Blocks clearly fixate on morphological features of the shark, such as its snout,  
 430 fins and tail. Figure 7 shows further examples of Phasor Blocks localising key characteristics of  
 431 different object categories, such as dog ears, elephant tusks and car wheels. Similar results visu-  
 alising activation maps of PsychoDW, and comparing them to those of ConvNeXt-S, are presented

in Figures A.5 (b) and A.6. Since KPCA-CAM only uses activations of the visualised layer and is uninfluenced by model predictions (e.g. via backpropagation in gradient-based CAMs), these results indicate that Phasor Blocks specialize to extract meaningful semantic object parts. This organisation is likely shaped by the presence of SVC downstream, which we show below encodes these parts into higher-level abstractions that support interpretable semantic reasoning. This is further support by the salience maps in Figure A.10 that shows although SVC alone does promote object-part focus, clear part-level isolation only emerges with the addition of Phasor Blocks and their complex features.

Interestingly, it appears that the imaginary components of Phasor Block activations capture more global features than the real components (i.e. a dog’s face vs. its ears). An initial clustering visualisation of Phasor Block activations is also presented in Figure A.8, which finds that observable clustering emerges both components which becomes increasingly pronounced at deeper layers. These findings suggest that Phasor Blocks learn a rich utilisation of complex-valued representations that is iteratively refined through each block.

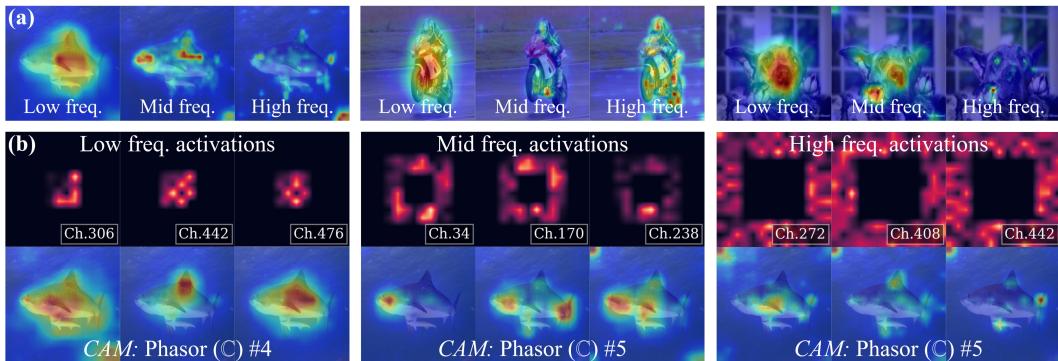


Figure 8: Psycho-B Phasor Block salience maps (via HiResCAM) conditioned on gradients (a) from individual Spectral Branch sub-bands and (b) from individual frequency domain feature channels.

For initial exploration of SVC’s encoding mechanisms, we use HiResCAM (Draelos & Carin, 2020) for gradient-based activation visualisation. It produces salience maps by element-wise multiplying layer activations with gradients backpropogated from model predictions, so in classification the salience regions have a high contribution to the class prediction. We extend this approach to isolate regions used by specific parts of SVC by first masking (setting to zero) gradients from the other components, enabling exploration of how SVC encodes Phasor Blocks features. First, we examine each of the three sub-bands created by PsychoNet’s Spectral Branches. After masking gradients of Hadamard Blocks for all but one of the sub-bands, Phasor Blocks’ salience regions reveal that SVC distributes object parts by scale. Figure 8 (a) shows that the low-frequency sub-band focuses on subjects broadly, while mid-high frequencies isolate more specific parts of different sizes. This aligns with frequency domain theory, in which low frequencies capture coarse spatial structure and higher frequencies finer detail and edges, supporting the view that SVC performs structured filtering in the frequency domain. We also isolate activations from individual Hadamard Block channels, showing that within each band, channels specialise to distinct object parts and correspond to distinct sparse frequency selections (Figure 8 (b)). This analysis was also applied to PsychoDW with similar results, shown in Figure A.7.

Overall, these result suggest that SVC learns a semantic intermediate representations that encodes selections of object parts. Given that SVC is placed immediately before the decision making (classification) layers of PsychoNet, it is likely selecting those most relevant to the task. In doing so, SVC functions as an abstraction bridging part extraction in Phasor Blocks and higher-level semantic reasoning, mirroring the role of abstractions used in psychovisual processing.

**Limitations and Future Work.** A key limitation of our work is that, although we show SVC organises and encodes selections of meaningful object parts, future research is still required to determine how the deeper semantic meaning of these abstractions should be interpreted in relation to broader notions of reasoning, such as those studied in neuroscience (Quiroga et al., 2005; Kriegeskorte et al., 2008). We will also explore addressing the high FLOP usage of PsychoNet by exploring

486 optimisations for complex-valued operations (e.g. employing Cauchy-Riemann identities (Ahlfors,  
487 1979)) as well as more sophisticated formulations of Phasor Blocks.

488 Additionally, it would also be insightful to explore applying PsychoNet to broader task types, par-  
489 ticularly image-to-image tasks like segmentation which may allow SVC to utilize wider frequency  
490 ranges than classification. Another focus will be practical applications in MRI, to explore how  
491 SVC’s semantic abstractions can enhance transparency and trustworthiness in models operating di-  
492 rectly on k-space data. Finally, it is also known that aliasing can afflict standard CNN architectures  
493 (Grabinski et al., 2022); future work should assess its impact on our frequency-domain representa-  
494 tions and whether mitigation can improve results.

495

## 496 5 CONCLUSION

497

498 In this work, we introduced Semantic Visual Coding (SVC), the first high-level vision representa-  
499 tion learnt in the frequency domain that produces sparse, data-driven coronal selections of discrete  
500 Fourier space. Our PsychoNet framework integrating SVCs show that it can maintain performance  
501 across multiple classification datasets, but is less depth-dependent, suggesting that SVC improves  
502 high-level processing previously done by deep spatial layers. In contrast to the unorganised pro-  
503 cessing of conventional CNNs, we find that PsychoNet clearly separates processing stages: Phasor  
504 Blocks extract semantically meaningful object parts, while SVCs encode and organise these parts  
505 into sparse, frequency domain representations used to make classification decisions that can be vi-  
506 sualized. This pipeline provides strong evidence that it may mimic intermediate abstractions used  
507 by the brain to separate feature extraction from higher cognition as suggested in previous neuro-  
508 science studies. While further work is required to understand the reasoning mechanisms of SVCs,  
509 it is clear that frequency domain abstraction is a promising direction for interpretable human-like  
510 model reasoning.

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   **ACKNOWLEDGEMENTS**  
541542   In accordance to ICLR 2026 guidelines, we acknowledge the use of large language models (LLMs)  
543   in preparing this manuscript. Their role was limited to assisting with editing and polishing writing.  
544545   Additional acknowledgements will be added after deanonymization.  
546547   **ETHICS STATEMENT**  
548549   All authors have reviewed the ICLR 2026 code of ethics and verified to the best of our knowledge  
550   that the work in our paper conforms with it. In particular, this work introduces a new theoretical  
551   framework, so it is unlikely to cause direct harm or negative impacts to society. Additionally, we  
552   only use open datasets such as ImageNet, so privacy concerns do not arise, though we acknowledge  
553   that these datasets contain known biases that may influence model behaviour.  
554555   **REPRODUCIBILITY STATEMENT**  
556557   We are committed to ensuring the reproducibility of our results. All experiments in this work were  
558   conducted on publicly available datasets, which have been appropriately cited. Detailed training  
559   recipes and hardware details are presented in Appendix C and Appendix D.1. Detailed model con-  
560   figurations are presented in Appendix B. After deanonymization, we will also release our code  
561   repository including training scripts, model weights and instructions to reproduce all of our results.  
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 REFERENCES  
595

- 596 Lars V. Ahlfors. *Complex Analysis*. McGraw-Hill, New York, 3rd edition, 1979. ISBN 978-  
597 0070006577.
- 598 Irwan Bello, W. Fedus, Xianzhi Du, E. D. Cubuk, A. Srinivas, Tsung-Yi Lin, Jonathon Shlens, and  
599 Barret Zoph. Revisiting ResNets: Improved Training and Scaling Strategies. *Neural Information*  
600 *Processing Systems*, 2021.
- 601 Patrick Le Callet, Abdelhakim Saadane, and Dominique Barba. Interactions of chromatic compo-  
602 nents in the perceptual quantization of the achromatic component. In *Human Vision and Elec-  
603 tronic Imaging IV*, volume 3644, pp. 121–128. SPIE, May 1999. doi: 10.1117/12.348432.
- 604 E. J. Candés and D. L. Donoho. Ridgelets: a key to higher-dimensional intermittency? *Philosoph-  
605 ical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 357  
606 (1760):2495–2509, 1999. doi: 10.1098/rsta.1999.0444.
- 607 S. S. Chandra, N. Normand, A. Kingston, J. Guédon, and I. Svalbe. Robust Digital Image Re-  
608 construction via the Discrete Fourier Slice Theorem. *Signal Processing Letters, IEEE*, 21(6):  
609 682–686, June 2014. ISSN 1070-9908. doi: 10.1109/LSP.2014.2313341.
- 610 Shekhar S Chandra, Marlon Bran Lorenzana, Xinwen Liu, Siyu Liu, Steffen Bollmann, and Stu-  
611 art Crozier. Deep learning in magnetic resonance image reconstruction. *Journal of Med-  
612 ical Imaging and Radiation Oncology*, 65(5):564–577, 2021. doi: <https://doi.org/10.1111/1754-9485.13276>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/1754-9485.13276>.
- 613 Lu Chi, Borui Jiang, and Yadong Mu. Fast Fourier Convolution. In *Advances in  
614 Neural Information Processing Systems*, volume 33, pp. 4479–4488. Curran Associates,  
615 Inc., 2020. URL [https://papers.nips.cc/paper\\_files/paper/2020/hash/2fd5d41ec6cfab47e32164d5624269b1-Abstract.html](https://papers.nips.cc/paper_files/paper/2020/hash/2fd5d41ec6cfab47e32164d5624269b1-Abstract.html).
- 616 Elizabeth Cole, Joseph Cheng, John Pauly, and Shreyas Vasanawala. Analysis of deep complex-  
617 valued convolutional neural networks for mri reconstruction and phase-focused applications.  
618 *Magnetic Resonance in Medicine*, 86(2):1093–1109, 2021. doi: <https://doi.org/10.1002/mrm.28733>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.28733>.
- 619 J. W Cooley, P. Lewis, and P. Welch. The finite Fourier transform. *Audio and Electroacoustics,  
620 IEEE Transactions on*, 17(2):77–85, June 1969. ISSN 0018-9278.
- 621 Muneer Dedmari, Sailesh Conjeti, Santiago Estrada, Phillip Ehses, Tony Stocker, and Martin Reuter.  
622 Complex fully convolutional neural networks for mr image reconstruction. In *Machine Learn-  
623 ing for Medical Image Reconstruction : first International Workshop, MLMIR 2018*, volume 1.  
624 Springer, 2018.
- 625 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. ImageNet: A large-scale hier-  
626 archical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*,  
627 pp. 248–255, June 2009. doi: 10.1109/cvpr.2009.5206848.
- 628 M.N. Do and M. Vetterli. The contourlet transform: an efficient directional multiresolution image  
629 representation. *IEEE Transactions on Image Processing*, 14(12):2091–2106, December 2005.  
630 ISSN 1941-0042. doi: 10.1109/TIP.2005.859376. URL <https://ieeexplore.ieee.org/document/1532309>.
- 631 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas  
632 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszko-  
633 reit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at  
634 scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event,  
635 Austria, May 3-7, 2021*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>.

- 648 Rachel Lea Draelos and Lawrence Carin. Use HiResCAM instead of Grad-CAM for faithful ex-  
 649 planations of convolutional neural networks. *arXiv preprint arXiv:2011.08891*, November 2020.  
 650 URL <https://arxiv.org/abs/2011.08891>.  
 651
- 652 Peng Gao, Jiasen Lu, Hongsheng Li, Roozbeh Mottaghi, and Aniruddha Kembhavi Kembhavi. Con-  
 653 tainer: Context Aggregation Network. *Neural Information Processing Systems*, 2021.  
 654
- 655 Rafael C. Gonzalez and Richard E. Woods. *Digital Image Processing 3rd Edition*. Prentice Hall,  
 656 January 2014.  
 657
- 658 Julia Grabinski, Steffen Jung, Janis Keuper, and Margret Keuper. Frequencylowcut pooling - plug  
 659 and play against catastrophic overfitting. In *Computer Vision – ECCV 2022: 17th European*  
 660 *Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIV*, pp. 36–57, Berlin,  
 661 Heidelberg, 2022. Springer-Verlag. ISBN 978-3-031-19780-2. doi: 10.1007/978-3-031-19781-9-  
 662 3. URL [https://doi.org/10.1007/978-3-031-19781-9\\_3](https://doi.org/10.1007/978-3-031-19781-9_3).  
 663
- 663 Julia Grabinski, Janis Keuper, and Margret Keuper. As large as it gets: Learning infinitely large  
 664 filters via neural implicit functions in the fourier domain. *ArXiv*, abs/2307.10001, 2023. URL  
 665 <https://api.semanticscholar.org/CorpusID:259982481>.  
 666
- 667 Bochen Guan, Jinnian Zhang, William A. Sethares, Richard Kijowski, and Fang Liu. Spectral  
 668 Domain Convolutional Neural Network. In *ICASSP 2021 - 2021 IEEE International Conference*  
 669 *on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2795–2799, June 2021. doi: 10.1109/  
 670 ICASSP39728.2021.9413409. ISSN: 2379-190X.  
 671
- 671 Jeanpierre V. Guedon, Dominique Barba, and Nicole Burger. Psychovisual image coding via an  
 672 exact discrete Radon transform. volume 2501, pp. 562–572, 1995. doi: 10.1117/12.206765.  
 673
- 674 Jean-Pierre V. Guédon, Nicolas Normand, Pierre Verbert, Benoit Parrein, and Florent Autrusseau.  
 675 Load-balancing and scalable multimedia distribution using the Mojette transform. *Internet Mul-  
 676 timedia Management Systems II*, 4519(1):226–234, 2001.  
 677
- 678 Jeanpierre Guédon. *The Mojette Transform: Theory and Applications*. John Wiley & Sons, March  
 679 2013. ISBN 9781118622933.  
 680
- 681 Jose Hanen and Dominique Barba. High-quality subband image coding of TV signals at 5 Mbit/s  
 682 with motion compensation interpolation and visually optimized scalar quantization. In *Visual*  
 683 *Communications and Image Processing '93*, volume 2094, pp. 1477–1485. SPIE, October 1993.  
 684 doi: 10.1117/12.157907.  
 685
- 685 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image  
 686 Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*,  
 687 pp. 770–778, 2016. doi: 10.1109/CVPR.2016.90.  
 688
- 689 Wen Hou and Cishen Zhang. Parallel-Beam CT Reconstruction Based on Mojette Transform and  
 690 Compressed Sensing. *International Journal of Computer and Electrical Engineering*, pp. 83–87,  
 691 2013. ISSN 17938163. doi: 10.7763/IJCEE.2013.V5.669.  
 692
- 692 Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q Weinberger. Densely connected  
 693 convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern*  
 694 *Recognition*, 2017.  
 695
- 696 Zhipeng Huang, Zhizheng Zhang, Cuiling Lan, Zheng-Jun Zha, Yan Lu, and Baining Guo. Adaptive  
 697 frequency filters as efficient global token mixers. In *2023 IEEE/CVF International Conference*  
 698 *on Computer Vision (ICCV)*, pp. 6026–6036, 2023. doi: 10.1109/ICCV51070.2023.00556.  
 699
- 700 D. H. Hubel and T. N. Wiesel. Receptive fields, binocular interaction and functional architecture in  
 701 the cat's visual cortex. *The Journal of Physiology*, 160(1):106–154.2, January 1962. ISSN 0022-  
 3751. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1359523/>.

- 702 Samira Kabri, Tim Roith, Daniel Tenbrinck, and Martin Burger. Resolution-invariant image clas-  
 703 sification based on fourier neural operators. In *Scale Space and Variational Methods in Com-*  
 704 *puter Vision: 9th International Conference, SSVM 2023, Santa Margherita Di Pula, Italy, May*  
 705 *21–25, 2023, Proceedings*, pp. 236–249, Berlin, Heidelberg, 2023. Springer-Verlag. ISBN 978-  
 706 3-031-31974-7. doi: 10.1007/978-3-031-31975-4\_18. URL [https://doi.org/10.1007/978-3-031-31975-4\\_18](https://doi.org/10.1007/978-3-031-31975-4_18).
- 707
- 708 Sachin Karmani, Thanushon Sivakaran, Gaurav Prasad, Mehmet Ali, Wenbo Yang, and Sheyang  
 709 Tang. KPCA-CAM: Visual Explainability of Deep Computer Vision Models Using Kernel PCA.  
 710 *IEEE International Workshop on Multimedia Signal Processing*, 2024. doi: 10.1109/mmfp61759.  
 711 2024.10743968.
- 712
- 713 Andrew Kingston and Imants Svalbe. Generalised finite radon transform for  $N \times N$  images. *Image*  
 714 *and Vision Computing*, 25(10):1620–1630, October 2007. ISSN 0262-8856. doi: 10.1016/j.imavis.2006.03.002.
- 715
- 716 Nikolaus Kriegeskorte, Marieke Mur, Douglas A. Ruff, Roozbeh Kiani, Jerzy Bodurka, Hossein  
 717 Esteky, Keiji Tanaka, and Peter A. Bandettini. Matching categorical object representations in  
 718 inferior temporal cortex of man and monkey. *Neuron*, 2008. doi: 10.1016/j.neuron.2008.10.043.
- 719
- 720 Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009. URL <https://api.semanticscholar.org/CorpusID:18268744>.
- 721
- 722 Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. ImageNet classification with deep  
 723 convolutional neural networks. *Communications of The ACM*, 60(6):84–90, May 2017. doi:  
 724 10.1145/3065386.
- 725
- 726 Lynn Le, Paolo Papale, K. Seeliger, Antonio Lozano, Thirza Dado, Feng Wang, Pieter R. Roelfsema,  
 727 M. van Gerven, Yağmur Güçlütürk, and Umut Güçlü. Monkeysee: Space-time-resolved recon-  
 728 structions of natural images from macaque multi-unit activity. *Neural Information Processing*  
 729 *Systems*, 2024.
- 730
- 731 ChiYan Lee, Hideyuki Hasegawa, and Shangce Gao. Complex-Valued Neural Networks: A Com-  
 732 prehensive Survey. *IEEE/CAA Journal of Automatica Sinica*, 9(8):1406–1426, August 2022. doi:  
 733 10.1109/jas.2022.105743.
- 734
- 735 J. Lee-Thorp, J. Ainslie, Ilya Eckstein, and Santiago Ontañón. FNet: Mixing Tokens with Fourier  
 736 Transforms. *North American Chapter of the Association for Computational Linguistics*, 2021.  
 737 doi: 10.18653/v1/2022.naacl-main.319.
- 738
- 739 Shaohua Li, Kaiping Xue, Bin Zhu, Chenkai Ding, Xindi Gao, David Wei, and Tao Wan. Falcon:  
 740 A fourier transform based approach for fast and secure convolutional neural network predictions.  
 741 In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8702–  
 742 8711, 2020a. doi: 10.1109/CVPR42600.2020.00873.
- 743
- 744 Wenhao Li, Wenqing Xie, and Zhifang Wang. Complex-valued densely connected convolutional  
 745 networks. In Jianchao Zeng, Weipeng Jing, Xianhua Song, and Zeguang Lu (eds.), *Data Science*,  
 746 pp. 299–309, Singapore, 2020b. Springer Singapore. ISBN 978-981-15-7981-3.
- 747
- 748 Shiqi Lin, Zhizheng Zhang, Zhipeng Huang, Yan Lu, Cuiling Lan, Peng Chu, Quanzeng You, Jiang  
 749 Wang, Zicheng Liu, Amey Parulkar, Viraj Navkal, and Zhibo Chen. Deep Frequency Filtering for  
 750 Domain Generalization. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recog-*  
 751 *nition (CVPR)*, pp. 11797–11807, Los Alamitos, CA, USA, June 2023. IEEE Computer Society.  
 752 doi: 10.1109/CVPR52729.2023.01135. URL <https://doi.ieee.org/10.1109/CVPR52729.2023.01135>.
- 753
- 754 Mengkun Liu, Licheng Jiao, Xu Liu, Lingling Li, Fang Liu, and Shuyuan Yang. C-CNN: Contourlet  
 755 Convolutional Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems*,  
 32(6):2636–2649, June 2021. ISSN 2162-2388. doi: 10.1109/TNNLS.2020.3007412. URL  
<https://ieeexplore.ieee.org/document/9145825>. Conference Name: IEEE  
 Transactions on Neural Networks and Learning Systems.

- 756 Mengkun Liu, Licheng Jiao, Xu Liu, Lingling Li, Fang Liu, Shuyuan Yang, and Xiangrong Zhang.  
 757 Bio-Inspired Multi-scale Contourlet Attention Networks. *IEEE transactions on multimedia*, pp.  
 758 1–16, January 2023. doi: 10.1109/tmm.2023.3304448.
- 759 Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining  
 760 Xie. A convnet for the 2020s. In *2022 IEEE/CVF Conference on Computer Vision and Pattern  
 761 Recognition (CVPR)*, pp. 11966–11976, 2022. doi: 10.1109/CVPR52688.2022.01167.
- 762 Nicolas Normand, Jean-Pierre Guédon, O. Philippe, and D. Barba. Controlled redundancy for image  
 763 coding and high-speed transmission. *Proc. of the SPIE - The International Society for Optical  
 764 Engineering*, 2727:1070–1081, 1996. ISSN 0277-786X.
- 765 Bruno A. Olshausen and David J. Field. Emergence of simple-cell receptive field properties by  
 766 learning a sparse code for natural images. *Nature*, 381(6583):607, June 1996. ISSN 1476-4687.  
 767 doi: 10.1038/381607a0. URL <https://www.nature.com/articles/381607a0>.
- 768 Benoit Parrein, Pierre Verbert, Nicolas Normand, and Jean-Pierre V. Guédon. Scalable multiple  
 769 descriptions on packets networks via the n-dimensional Mojette transform. *Quality of Service  
 770 over Next-Generation Data Networks*, 4524(1):243–252, 2001.
- 771 Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito,  
 772 Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in  
 773 pytorch. In *NIPS-W*, 2017.
- 774 Rodrigo Quiroga. Concept cells: the building blocks of declarative memory functions. *Nature  
 775 Reviews Neuroscience*, 2012. doi: 10.1038/nrn3251.
- 776 Rodrigo Quiroga, Leila Reddy, Gabriel Kreiman, Christof Koch, and Itzhak Fried. In-  
 777 variant visual representation by single neurons in the human brain. *Nature*, 2005. doi:  
 778 10.1038/nature03687.
- 779 Yongming Rao, Wenliang Zhao, Zheng Zhu, Jie Zhou, and Jiwen Lu. Gfnet: Global filter networks  
 780 for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9):  
 781 10960–10973, 2023. doi: 10.1109/TPAMI.2023.3263824.
- 782 Dario L. Ringach. Spatial structure and symmetry of simple-cell receptive fields in macaque primary  
 783 visual cortex. *Journal of Neurophysiology*, 88(1):455–463, July 2002. doi: 10.1152/jn.2002.88.  
 784 1.455.
- 785 Oren Rippel, Jasper Snoek, and Ryan P. Adams. Spectral Representations for Convolutional Neural  
 786 Networks. In *Proceedings of the 28th International Conference on Neural Information Processing  
 787 Systems - Volume 2*, NIPS’15, pp. 2449–2457, Cambridge, MA, USA, 2015. MIT Press.
- 788 A. Saadane, H. Sénane, and D. Barba. Visual Coding: Design of Psychovisual Quantizers. *Journal  
 789 of Visual Communication and Image Representation*, 9(4):381–391, December 1998. ISSN 1047-  
 790 3203. doi: 10.1006/jvci.1998.0393.
- 791 Abdelhakim Saadane, Hakim Senane, and Dominique Barba. Design of psychovisual quantizers for  
 792 a visual subband image coding. In *Visual Communications and Image Processing ’94*, volume  
 793 2308, pp. 1446–1453. SPIE, September 1994. doi: 10.1117/12.185903.
- 794 Abdelhakim Saadane, Nachida Bekkat, and Dominique Barba. On the masking effects in a per-  
 795 ceptually based image quality metric. In *Imaging and vision systems: theory, assessment and  
 796 applications*, pp. 161–177. Nova Science Publishers, Inc., USA, January 2001. ISBN 978-1-  
 797 59033-033-3.
- 798 Simone Scardapane, Steven Van Vaerenbergh, Amir Hussain, and Aurelio Uncini. Complex-valued  
 799 Neural Networks with Non-parametric Activation Functions, February 2018. URL <http://arxiv.org/abs/1802.08026> [cs].
- 800 H. Senane, A. Saadane, and D. Barba. Image coding in the context of a psychovisual image repres-  
 801 entation with vector quantization. In *Proceedings., International Conference on Image Processing*,  
 802 volume 1, pp. 97–100 vol.1, October 1995. doi: 10.1109/ICIP.1995.529048.

- 810 Jean-Luc Starck, E.J. Candes, and D.L. Donoho. The curvelet transform for image denoising. *IEEE*  
 811 *Transactions on Image Processing*, 11(6):670–684, June 2002. ISSN 1941-0042. doi: 10.1109/  
 812 TIP.2002.1014998. Conference Name: IEEE Transactions on Image Processing.
- 813
- 814 I. Tolstikhin, N. Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner,  
 815 Jessica Yung, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. MLP-  
 816 Mixer: An all-MLP Architecture for Vision. In *Neural Information Processing Systems*, 2021.
- 817
- 818 Chiheb Trabelsi, Olexa Bilaniuk, Ying Zhang, Dmitriy Serdyuk, Sandeep Subramanian, Joao Felipe  
 819 Santos, Soroush Mehri, Negar Rostamzadeh, Yoshua Bengio, and Christopher J. Pal. Deep Com-  
 820 plex Networks. In *International Conference on Learning Representations*, February 2018. URL  
 821 <https://openreview.net/forum?id=H1T2hmZAb>.
- 822
- 823 Bhavya Vasudeva, Puneesh Deora, Saumik Bhattacharya, and Pyari Mohan Pradhan. Compressed  
 824 sensing mri reconstruction with co-vegan: Complex-valued generative adversarial network. In  
 825 *2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 1779–1788,  
 826 2022. doi: 10.1109/WACV51458.2022.00184.
- 827
- 828 Pierre Verbert, Jean-Pierre V. Guédon, and Benoit Parrein. Distributed and compressed multimedia  
 829 transmission using a discrete backprojection operator. *Internet Multimedia Management Systems*  
 830 *III*, 4862(1):315–325, 2002. doi: 10.1111/12.473047.
- 831
- 832 Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual trans-  
 833 formations for deep neural networks. In *2017 IEEE Conference on Computer Vision and Pattern*  
 834 *Recognition (CVPR)*, pp. 5987–5995, 2017. doi: 10.1109/CVPR.2017.634.
- 835
- 836 Saurabh Yadav and Koteswar Rao Jerripothula. FCCNs: Fully Complex-valued Convolutional  
 837 Networks using Complex-valued Color Model and Loss Function. In *2023 IEEE/CVF In-  
 838 ternational Conference on Computer Vision (ICCV)*, pp. 10655–10664, October 2023. doi:  
 839 10.1109/ICCV51070.2023.00981. URL <https://ieeexplore.ieee.org/document/10377516>. ISSN: 2380-7504.
- 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 Appendices

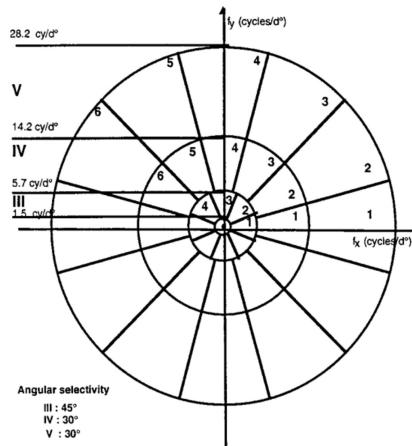
865  
 866  
 867 In the following we present appendices to our work, structured as follows: Appendix A presents  
 868 additional background material. Appendix B provides detailed information about the architectural  
 869 configurations of all models used in our experiments. Appendix C presents results, dataset informa-  
 870 tion and training recipes for all of our classification experiments. Appendix D.1 provides full results  
 871 and details for PsychoNet architectural ablation studies.

## 872 A BACKGROUND

873 In this section we present additional background and details about psychovisual coding, the Fourier  
 874 Transform and complex-valued networks.

### 875 A.1 PSYCHOVISUAL CODING

876 Our work is inspired by groundbreaking research conducted in the 1990s by French researchers led  
 877 by Dominique Barba in understanding the human aspect of mammalian vision, i.e. the psychovi-  
 878 sual capability of the human brain for visual perception arising from the need for early television  
 879 signal compression (Hanen & Barba, 1993). At the time, statistical approaches based on Shannon’s  
 880 information theory and rectilinear methods such as discrete (Haar) wavelets were popular, and they  
 881 argued that these approaches were sub-optimal because they treated all errors equally. They would  
 882 propose psychovisual quantizers as an efficient form of image coding that would retain the impor-  
 883 tant image information pertaining to its interpretation by the human vision system and quantization  
 884 matched the detection thresholds of the visual cortex (Senane et al., 1995). These quantizers were  
 885 proposed to be the coronas of the 2D Fourier space, where the model of the vision system assumes  
 886 Fourier space is analyzed using radial symmetric functions (Saadane et al., 1994; 1998) (see Fig-  
 887 ure A.1), which they showed can also be mapped to colors in human vision (Callet et al., 1999). The  
 888 premise is that visual recognition and feature extraction could be performed by selecting coronal  
 889 sectors of Fourier space directly through the quantisation of adjacent frequencies, thereby providing  
 890 directional band limited filtering within the scene. Their psychophysical experiments were also  
 891 used to select the optimal sub-bands that allowed image compression that was difficult for humans  
 892 to distinguish (Saadane et al., 2001). Our SVC is a data-driven adaptation of this approach, using  
 893 band-limited frequency filters to learn sparse frequency selections using supervisory signals from a  
 894 classification task.



913 Figure A.1: Coronal frequency sub-bands used in psychovisual coding from Saadane et al. (1998).  
 914

915  
 916 This work on visual codes over the course of a decade would result in among the first uses of  
 917 vector quantization for image coding (Senane et al., 1995), a perceptually based image quality met-  
 ric (Saadane et al., 2001) and one of the foundations of discrete projection theory, where a central

slice theorem is established for discrete Fourier space based as exact 1D forms of these psychovisual radial functions as slices and therefore projections in image space (Guedon et al., 1995). This work would even pioneer the use of the wavelet transform to projection data before it would be formalized as ridgelets by Candes and Donoho (Candés & Donoho, 1999). The Mojette transform would itself form the basis of an entire area of discrete tomography that creates discrete projections of images (Normand et al., 1996) in diverse areas such as image reconstruction (Kingston & Svalbe, 2007; Chandra et al., 2014) and compression (Guédon et al., 2001), computed tomography (Hou & Zhang, 2013) and network transmission (Parrein et al., 2001; Verbert et al., 2002). Although the number of publications is too numerous to list here, a summary of these works and areas can be found in the Mojette transform book (Guédon, 2013).

## A.2 THE FOURIER TRANSFORM AND THE FREQUENCY DOMAIN

In Section 3 we only describe the 2D DFT as digital images are discrete 2D signals in the spatial domain, while the standard FT operates on continuous signals. The DFT is derived by first viewing a discrete signal as the product of a continuous signal and a sequence of unit impulses (sampling), applying the FT to yield a continuous function in the frequency domain, then sampling it again to discretize it. Detailed derivations of both the FT and DFT may be found in most image processing texts, such as (Gonzalez & Woods, 2014). There are also inverse transforms, namely the Inverse Fourier Transform (IFT) and Inverse Discrete Fourier Transform (IDFT), for transforming frequency domain signals back into the spatial domain. While we do not use them in PsychoNet, they reflect the duality between the spatial and frequency domains - any operation in one domain has a counterpart in the other. The most famous example of this relationship is the Convolution Theorem.

Let  $x[u, v], y[u, v], u, v \in 0, \dots, N - 1$  be two discrete  $N \times N$  spatial signals. The circular convolution of these two signals is defined as:

$$x[u, v] * y[u, v] = \frac{1}{N^2} \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} x[m, n] y[((u - m))_N, ((v - n))_N] \quad (2)$$

where  $(.)_N$  denotes modulo  $N$ . The Convolution Theorem (Gonzalez & Woods, 2014) then states that:

$$\mathcal{F}[x * y] = \mathcal{F}[x] \odot \mathcal{F}[y] \text{ or equivalently } x * y = \mathcal{F}^{-1}[\mathcal{F}[x] \odot \mathcal{F}[y]] \quad (3)$$

where  $\mathcal{F}[\cdot]$  and  $\mathcal{F}^{-1}[\cdot]$  denote the DFT and IDFT, and  $\odot$  the Hadamard product. Hence, circular convolution in the spatial domain is equivalent to applying the Hadamard product in the frequency domain. As such, the frequency domain is highly conducive to global representations, since each element of an image's frequency spectra presents a unique global view of the image, analogous to convolving it with a directional striped kernel.

In practice, the DFT and IDFT are computed using the Fast Fourier Transform and Inverse Fast Fourier Transform respectively (Cooley et al., 1969). Note that if  $x$  and  $y$  were multi-channel features instead, i.e. of dimension  $d \times N \times N$  for  $d$  channels like the input features and learnt filters of our Hadamard Blocks, then the frequency domain Hadamard product is equivalent to circular *depthwise* convolution in the spatial domain. Unlike the Conv2D operation of CNNs, this does not mix channels, which is why both of our Hadamard Blocks and the Global Filter block from Rao et al. (2023) include explicit channel-mixing via  $1 \times 1$  convolution layers.

## A.3 COMPLEX-VALUED NEURAL NETWORKS

Most work for complex-valued neural networks involve developing components of these networks to work in the complex domain, such as activation functions (Scardapane et al., 2018). Most complex-valued CNNs use the network blocks introduced by Trabelsi et al. (2018). The distributive property of convolution allows convolution between a complex input  $\mathbf{h} = \mathbf{a} + i\mathbf{b}$  and a complex kernel  $\mathbf{W} = \mathbf{W}_R + i\mathbf{W}_I$  to be decomposed into four real-valued component wise convolutions:

$$\mathbf{W} * \mathbf{h} = (\mathbf{W}_R * \mathbf{a} - \mathbf{W}_I * \mathbf{b}) + i(\mathbf{W}_I * \mathbf{a} + \mathbf{W}_R * \mathbf{b}) \quad (4)$$

Consequently, complex-valued convolution layers are usually more computationally and memory intensive (additionally stores imaginary features) than real-valued ones. Trabelsi et al. (2018) also

972 developed complex normalization methods and activation functions. Complex-valued modules in  
973 PsychoNet use the complex-valued convolution (CConv) and batch-normalization (CBN) layers  
974 from Trabelsi et al. (2018), and a naïve adaptation of the GELU activation function (CGELU) which  
975 just applies the original function to real and imaginary channels separately.

976 When applying complex-valued networks to real-valued images, most works use a small initial  
977 module to convert the input into complex-valued features. However, such approaches have yielded  
978 only minor improvements in the past over directly using real-valued networks (Trabelsi et al., 2018;  
979 Li et al., 2020b). Accordingly, recent complex-valued networks predominantly focus on domains  
980 with naturally complex data, such as MRI, radar and audio signal processing (Dedmari et al., 2018;  
981 Vasudeva et al., 2022; Cole et al., 2021; Lee et al., 2022; Trabelsi et al., 2018). To try bridge this  
982 gap, a complex-valued colour space by reinterpreting the cylindrical coordinates of the HSV colour  
983 model as 2D magnitude and phase was developed (Yadav & Jerripothula, 2023). They applied this  
984 to standard complex-valued CNNs, improving results on common image classification tasks, but  
985 retained the high complexity of complex-valued networks. On the other hand, PsychoNet primarily  
986 uses real-valued modules that learn to generate *complementary* complex-valued features to given  
987 real features, as described in Section 3.

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 **B MODEL CONFIGURATIONS**  
10271028 Here we provide full details of the architectural configurations of all of our models, For all tables,  
1029 we use the ResNet approach of counting the number of model layers as the number of convolutional  
1030 and linear layers; each element-wise filter block in Hadamard Blocks are also counted as one layer.  
10311032 **B.1 CNN BASELINES**  
10331034 The ResNet50, 101 and 152 models we use are from He et al. (2016) and are implemented in most  
1035 common deep learning frameworks (we use the one from PyTorch (Paszke et al., 2017)). For  
1036 ResNet270, we follow the block configurations in Bello et al. (2021), but do not implement any of  
1037 the newer blocks/layers they also introduce, so it purely just adds more residual bottleneck blocks  
1038 (ResBlocks) to ResNet152 for fair scaling. Table A.1 compares the sizes of the four ResNet models  
1039 as well as their block configurations, grouped by feature resolution (which are  $56 \times 56$ ,  $28 \times 28$ ,  $14 \times$   
1040  $14$  and  $7 \times 7$ ).  
10411042 Table A.1: ResNet block configurations.  
1043

Model	Parameters (M)	# Layers	# Blocks
ResNet50	25.56	54	[3-4-6-3]
ResNet101	44.55	105	[3-4-23-3]
ResNet152	60.19	156	[3-8-36-3]
ResNet270	89.60	276	[4-29-53-3]

1044 For ConvNeXt-S, we follow the original implementation in Liu et al. (2022).  
10451051 **B.2 PSYCHONET**  
10521053 Figure A.2 summarises the key configuration details of each PsychoNet variant, namely the fea-  
1054 ture resolution and channel width at each Phasor Block and the number of filters used in the SVC  
1055 module. Further architectural details for each model are reported in Tables A.3 through A.7. For  
1056 Phasor Blocks, we list each layer using the ‘**resolution**: layer configuration’ format. The ResNet-  
1057 based PsychoNet models use the same initial input embedding layer as ResNet ( $7 \times 7$  Conv2D and  
1058 maxpooling) is used, while PsychoDW uses the same  $4 \times 4$  patch embeddings as ConvNeXt-S.  
1059 Interestingly, we found that using the initial layers of ResNet50, instead of ConvNeXt-S, in our  
1060 ConvNeXt-S based PsychoDW actually yielded better results (approx.  $\uparrow 0.5\%$  top-1 accuracy on  
1061 ImageNet-1K), so we chose to use it for the model. However, we do change all ConvBlocks in  
1062 Phasor ( $\mathbb{C}$ ) (see Figure A.2) to depthwise convolution blocks to maintain general faithfulness to the  
1063 ConvNeXt model.  
10641065 Finally, the companding operation we apply after taking the 2D FFT (in Figure 3) simply zeros the  
1066 DC component and applies the element-wise function:  
1067

1068 
$$x \in \mathbb{C}, \quad \text{Compan} : x \rightarrow |x|^{\frac{1}{1.25}} \cdot \exp(i\angle x) \quad (5)$$
  
1069 where  $|x|$  denotes the magnitude of  $x$  and  $\angle x$  its phase. Since the exponent applied to the magnitude  
1070 is  $\in (0, 1)$ , this function compresses frequencies of large magnitude (i.e. frequencies very close to  
the DC component), and expands the magnitude of those further from it.  
1071  
1072  
1073  
1074  
1075  
1076  
1077  
1078  
1079

1080  
 1081 Table A.2: Configuration summary of different PsychoNet variants. For Phasor Blocks, we display  
 1082 **resolution:** [#channels per block], and ‘(I)’ denotes a Phasor Block (I)). #Filters denotes the  
 1083 number of channels of element-wise filters per sub-band of SVC. #Layers show overall layers /  
 1084 complex convolution layer counts.

Model	Phasor Blocks	#Filters	#Layers	Params (M)
Psycho-S	<b>14 × 14:</b> [256 (I), 256, 384, 512, 512]	512	65 / 9	25.35
Psycho-B	<b>28 × 28:</b> [256 (I), 256, 256, 384] <b>14 × 14:</b> [384, 384, 512, 512, 512]	512	93 / 13	42.01
Psycho-L	<b>28 × 28:</b> [256 (I), 512, 512, 512] <b>14 × 14:</b> [512, 512, 512, 512]	512	93 / 13	61.28
Psycho-H	<b>28 × 28:</b> [256 (I), 512, 512, 512] <b>14 × 14:</b> [512, 512, 512, 640, 1024]	1024	93 / 13	88.61
Psycho-DW	<b>28 × 28:</b> [256 (I), 256, 256, 512] <b>14 × 14:</b> [512, 1024, 1024, 1024, 1024]	2048	109 / 13	49.512

1097  
 1098 Table A.3: Detailed architecture of Psycho-S.  
 1099

Psycho-S - based on ResNet50	
Parameters (M)	25.35
# Layers (overall)	65
# Layers (complex)	9
Blocks	
Input layer	Conv2D( $7 \times 7$ , $d_{in}=3$ , $d_{out}=64$ , stride=2), MaxPool( $3 \times 3$ , stride=2)
Initial CNN layers	First 7 ResBlocks from ResNet50 (first two resolution stages).
Phasor Blocks	
	<b>14 × 14:</b> (I) [ $d_{in}=128$ , $d_{out}=256$ , stride=2]
	<b>14 × 14:</b> (C) [ $d_{in}=256$ , $d_{out}=256$ ]
	<b>14 × 14:</b> (C) [ $d_{in}=256$ , $d_{out}=384$ ]
	<b>14 × 14:</b> (C) [ $d_{in}=384$ , $d_{out}=512$ ]
	<b>14 × 14:</b> (C) [ $d_{in}=512$ , $d_{out}=512$ ]
Spectral filters	Sub-bands ([crop, drop]): [14, 8], [8, 4], [4, 1], d_filter = 512
Output layer	Average pool, ComplexLinear( $d_{in}=1536$ , $d_{out}=1000$ ), Softmax

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

Table A.4: Detailed architecture of Psycho-B.

<b>Psycho-B architecture - based on ResNet101</b>	
Parameters (M)	42.01
# Layers (overall)	93
# Layers (complex)	13
<b>Blocks</b>	
Input layer	Conv2D( $7 \times 7$ , $d_{in}=3$ , $d_{out}=64$ , stride=2), MaxPool( $3 \times 3$ , stride=2)
Initial CNN layers	First 7 ResBlocks from ResNet101 (first two resolution stages).
<b>28 × 28: (I)</b> [ $d_{in}=128$ , $d_{out}=256$ ]	
<b>28 × 28: (C)</b> [ $d_{in}=256$ , $d_{out}=256$ ]	
<b>28 × 28: (C)</b> [ $d_{in}=256$ , $d_{out}=256$ ]	
<b>28 × 28: (C)</b> [ $d_{in}=256$ , $d_{out}=384$ ]	
Phasor Blocks	
<b>14 × 14: (C)</b> [ $d_{in}=384$ , $d_{out}=384$ , stride=2]	
<b>14 × 14: (C)</b> [ $d_{in}=384$ , $d_{out}=384$ ]	
<b>14 × 14: (C)</b> [ $d_{in}=384$ , $d_{out}=512$ ]	
<b>14 × 14: (C)</b> [ $d_{in}=512$ , $d_{out}=512$ ]	
<b>14 × 14: (C)</b> [ $d_{in}=512$ , $d_{out}=512$ ]	
Spectral filters	Sub-bands ([crop, drop]): [14, 8], [8, 4], [4, 1], d_filter = 512
Output layer	Average pool, ComplexLinear( $d_{in}=1536$ , $d_{out}=1000$ ), Softmax

Table A.5: Detailed architecture of Psycho-L.

<b>Psycho-L architecture - based on ResNet152</b>	
Parameters (M)	61.28
# Layers (overall)	93
# Layers (complex)	13
<b>Blocks</b>	
Input layer	Conv2D( $7 \times 7$ , $d_{in}=3$ , $d_{out}=64$ , stride=2), MaxPool( $3 \times 3$ , stride=2)
Initial CNN layers	First 7 ResBlocks from ResNet152.
<b>28 × 28: (I)</b> [ $d_{in}=128$ , $d_{out}=256$ ]	
<b>28 × 28: (C)</b> [ $d_{in}=256$ , $d_{out}=512$ ]	
<b>28 × 28: (C)</b> [ $d_{in}=512$ , $d_{out}=512$ ]	
<b>28 × 28: (C)</b> [ $d_{in}=512$ , $d_{out}=512$ ]	
Phasor Blocks	
<b>14 × 14: (C)</b> [ $d_{in}=512$ , $d_{out}=512$ , stride=2]	
<b>14 × 14: (C)</b> [ $d_{in}=512$ , $d_{out}=512$ ]	
<b>14 × 14: (C)</b> [ $d_{in}=512$ , $d_{out}=512$ ]	
<b>14 × 14: (C)</b> [ $d_{in}=512$ , $d_{out}=512$ ]	
<b>14 × 14: (C)</b> [ $d_{in}=512$ , $d_{out}=512$ ]	
Spectral filters	Sub-bands ([crop, drop]): [14, 8], [8, 4], [4, 1], d_filter = 512
Output layer	Average pool, ComplexLinear( $d_{in}=1536$ , $d_{out}=1000$ ), Softmax

1188

1189

1190

1191

Table A.6: Detailed architecture of Psycho-H.

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

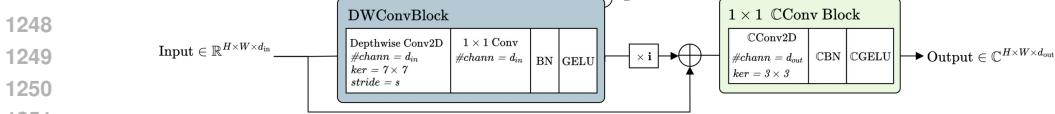
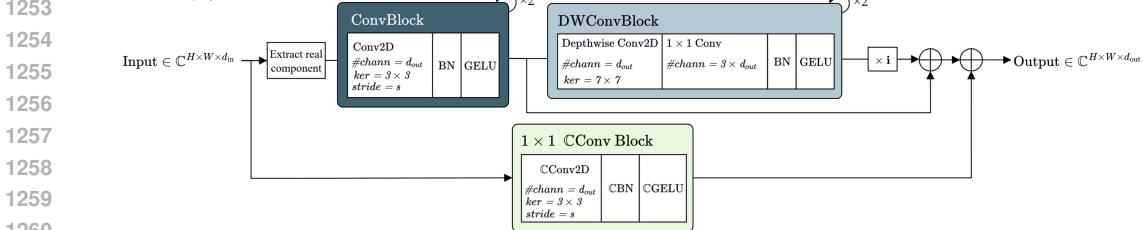
**Psycho-H architecture - based on ResNet270**

Parameters (M)	88.61
# Layers (overall)	93
# Layers (complex)	13
<b>Blocks</b>	
Input layer	Conv2D( $7 \times 7$ , $d_{in}=3$ , $d_{out}=64$ , stride=2), MaxPool( $3 \times 3$ , stride=2)
Initial CNN layers	First 7 ResBlocks from ResNet270.
<b>Phasor Blocks</b>	
	<b>28 × 28:</b> (I) [ $d_{in}=128$ , $d_{out}=256$ ]
	<b>28 × 28:</b> (C) [ $d_{in}=256$ , $d_{out}=512$ ]
	<b>28 × 28:</b> (C) [ $d_{in}=512$ , $d_{out}=512$ ]
	<b>28 × 28:</b> (C) [ $d_{in}=512$ , $d_{out}=512$ ]
<b>Spectral filters</b>	
	Sub-bands ([crop, drop]): [14, 8], [8, 4], [4, 1], d_filter = 1024
Output layer	Average pool, ComplexLinear( $d_{in}=3072$ , $d_{out}=1000$ ), Softmax

Table A.7: Detailed architecture of PsychoDW.

**PsychoDW architecture - based on ConvNeXt-S**

Parameters (M)	49.512
# Layers (overall)	109
# Layers (complex)	13
<b>Blocks</b>	
Input layer	Conv2D( $7 \times 7$ , $d_{in}=3$ , $d_{out}=64$ , stride=2), MaxPool( $3 \times 3$ , stride=2)
Initial CNN layers	First 7 ResBlocks from ResNet50.
<b>Phasor Blocks</b>	
	<b>28 × 28:</b> (I) [ $d_{in}=128$ , $d_{out}=256$ ]
	<b>28 × 28:</b> (C) [ $d_{in}=256$ , $d_{out}=256$ ]
	<b>28 × 28:</b> (C) [ $d_{in}=256$ , $d_{out}=256$ ]
	<b>28 × 28:</b> (C) [ $d_{in}=256$ , $d_{out}=512$ ]
<b>Spectral filters</b>	
	Sub-bands ([crop, drop]): [14, 8], [8, 4], [4, 1], d_filter = 1024
Output layer	Average pool, ComplexLinear( $d_{in}=3072$ , $d_{out}=1000$ ), Softmax

1242 B.3 PHASOR BLOCK ARCHITECTURE  
12431244 Figure A.2 provides detailed architectural diagrams of Phasor Blocks, with key design choices dis-  
1245 cussed below.  
12461247 **Phasor (I) Block**1252 **Phasor (C) Block**1253 Figure A.2: Further architecture details for the Phasor Blocks presented in Figure 4. For ConvNeXt-  
1254 based PsychoNet, we replace the two ConvBlocks at the start of Phasor (C) blocks with two DW-  
1255 ConvBlocks with the same number of channels. CConv/BN/GELU denote complex-valued conve-  
1256 lution, batch norm and GELU operations - see Appendix A.3. The following PsychoNet architecture  
1257 tables specify the values of  $d_{in}$ ,  $d_{out}$  and stride ( $s$ ) for all of their Phasor Blocks.  
1258

1259

1260 Phasor (I) blocks generate an initial set of imaginary components using depthwise convolution  
1261 ('DWConv') blocks, comprising pairs of depthwise and  $1 \times 1$  convolution layers. This configuration  
1262 decouples spatial and channel mixing, which is intended to encourage cross-channel interactions  
1263 without interfering with spatial relationships. In natural complex signals, the real and imaginary  
1264 components carry complementary information for the same spatial location (Gonzalez & Woods,  
1265 2014; Lee et al., 2022), so it is likely important that our generated imaginary features do not sig-  
1266 nificantly introduce new spatial information. A  $1 \times 1$  complex convolution block then mixes the  
1267 real and imaginary features. Subsequently, Phasor (C) blocks further refine the complex represen-  
1268 tations. The top branch generates new real and imaginary features, while the bottom channel-mixes  
1269 the original features and combines them with the new ones. For ConvNeXt-based PsychoNet, we  
1270 replace Phasor (C) 's regular convolution ('Conv') blocks with further DWConv blocks with  $7 \times 7$   
1271 kernel size, matching ConvNeXt's main computational block.  
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 

## C CLASSIFICATION EXPERIMENTS

1297

1298 In this section we present detailed results, dataset details and training recipes for all classification  
1299 experiments conducted.
1300

1301 

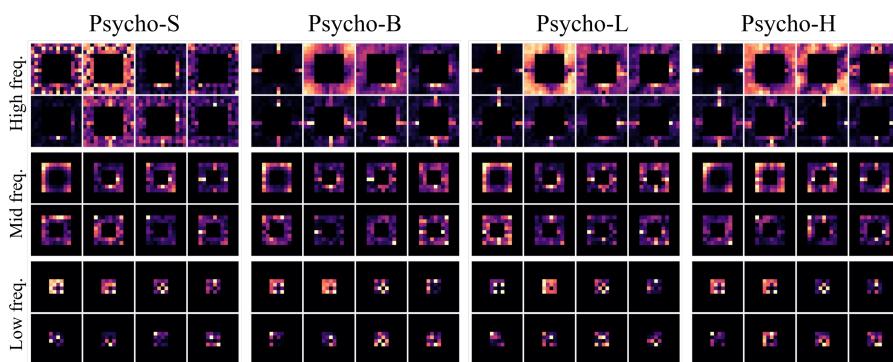
### C.1 IMAGENET-1K

1302

1303 We use the standard large ImageNet-1K subset from (Deng et al., 2009) containing  $\sim 1.2$  million  
1304 training and  $\sim 50000$  images for validation/testing. Table A.8 presents the training recipe used for  
1305 ImageNet experiments.
1306

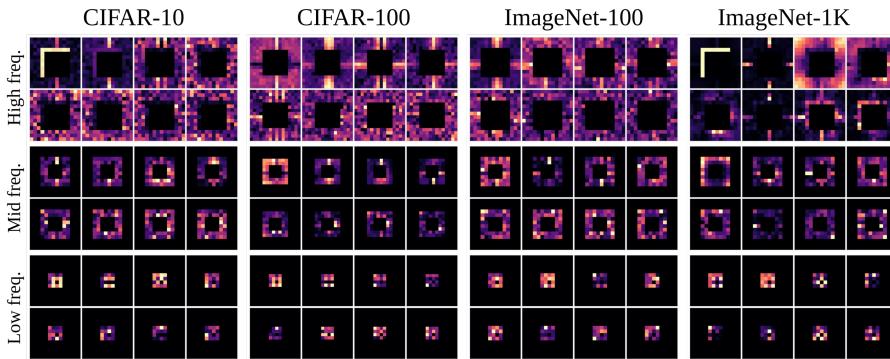
Table A.8: ImageNet training recipe

Setting	Value
Image size	$224 \times 224$
Epochs	90
Batch size (overall, not per GPU)	1024
Loss	Cross entropy
Optimizer	AdamW ( $\beta_1 = 0.9, \beta_2 = 0.999$ )
Scheduler	cosine
Initial learning rate (LR)	$5 \cdot 10^{-4}$
Warmup	warmup LR = $10^{-6}$ , 5 epochs
Learning rate decay	min. LR = $10^{-5}$ , 12 epochs
Augmentation	resize, crop, interpolate, horizontal flip, RandAugment, MixUp, CutMix, label smoothing
GPU	2x NVIDIA H100: Psycho-B, ResNet101, all ‘Big’ sized ablation models 2x AMD MI300X: Psycho-S, ResNet50 4x AMD MI300X: All other models

1325 Table A.8 presents all ImageNet-1K experiment results. PsychoNet moderately improves top-1  
1326 accuracy for all ResNet baselines ( $\uparrow 0.82\%, 0.41\%, 0.26\%$  and  $0.44\%$  vs. ResNet50 to 270), and in-  
1327 curs a small decrease for ConvNeXt-S ( $\downarrow 0.19$ ). Figure A.3 compares SVC filters learnt by different  
1328 ResNet-based PsychoNet sizes, showing that with larger model size, the filters become increasingly  
1329 structured and sparser, with clearer frequency selectivity and reduced noise. Figure A.4 compares  
1330 SVC filters learnt by Psycho-B on ImageNet-1K to the smaller resolution/size datasets in Appendix  
1331 C.2. It is evident that increasing image resolution and dataset size both yield much sparser filters.  
1332 These results suggest that the sparse patterns correspond to a data-driven representation naturally  
1333 emergent from visual information.
1334

1347 Figure A.3: Top principal components of SVC filters learnt by different sized ResNet-based Psy-  
1348 choNet models on ImageNet-1K. ‘High/mid/low freq.’ refer to the [14, 8], [8, 4] and [4, 1] frequency  
1349 sub-bands created by Spectral Branches.

1350  
1351 Table A.9: ImageNet-1K classification results. Each pair of rows (separated by horizontal lines)  
1352 compares a baseline CNN and the PsychoNet based on it. FLOPs were measured using a single  
1353  $224 \times 224$  input.

Model	Top-1 Acc. (%)	Top-5 Acc. (%)	Layers	Params (M)	FLOPs (G)	GPU
ResNet50	76.044	92.992	54	25.56	8.18	$2 \times$ MI300X
Psycho-S	76.864	93.386	65	25.35	12.31	$2 \times$ MI300X
ResNet101	78.428	94.220	105	44.55	15.60	$2 \times$ H100
Psycho-B	78.846	94.600	93	42.01	30.13	$2 \times$ H100
ResNet152	79.586	94.684	156	60.19	23.03	$4 \times$ MI300X
Psycho-L	79.848	95.056	93	61.28	54.47	$4 \times$ MI300X
ResNet270	80.012	95.088	276	89.60	40.50	$4 \times$ MI300X
Psycho-H	80.454	95.290	93	88.61	64.12	$4 \times$ MI300X
ConvNeXt-S	80.780	95.488	113	50.22	17.36	$2 \times$ MI300X
Psycho-DW	80.590	95.384	106	49.51	27.42	$2 \times$ MI300X



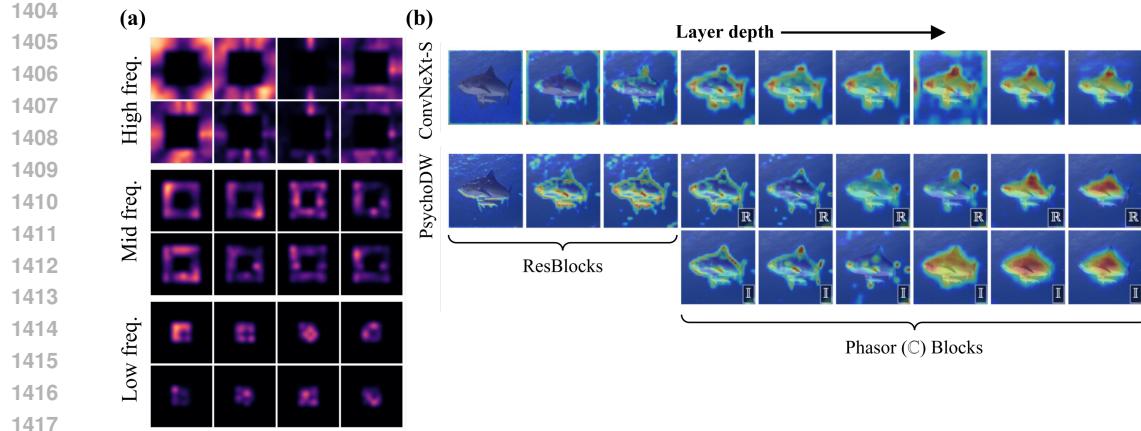
1368  
1369 Figure A.4: Top principal components of SVC filters learnt by Psycho-B on different resolution and  
1370 size datasets. ‘High/mid/low freq.’ refer to the [14, 8], [8, 4] and [4, 1] frequency sub-bands created  
1371 by Spectral Branches.

### C.1.1 PSYCHODW REPRESENTATION ANALYSIS

1380  
1381 Figures A.5 through A.7 present qualitative visualisations and analysis for PsychoDW identical to  
1382 those applied to Psycho-B we presented in Section 4. Overall, these show similar results:

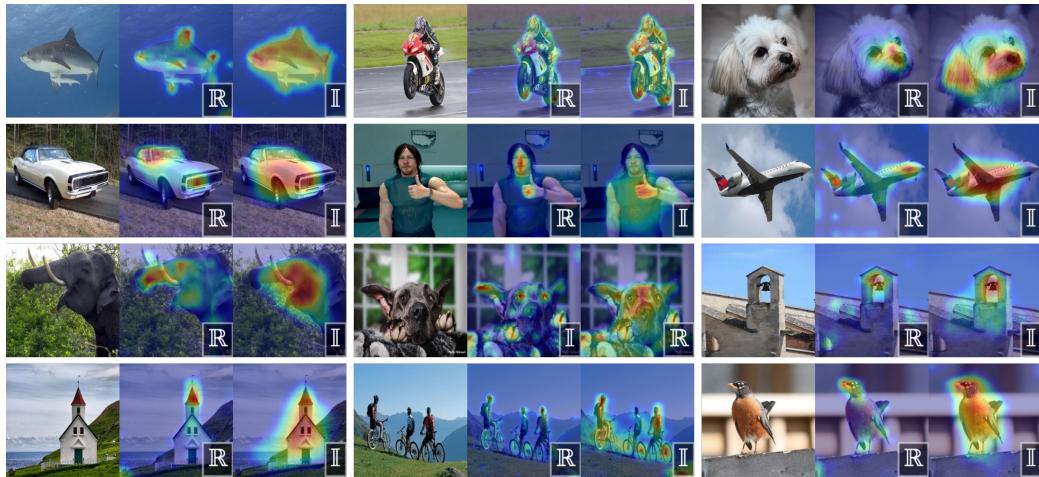
- 1383 • Figure A.5 **(a)** shows that PsychoDW’s SVC filters also learn sparse selections of frequencies across each sub-band.
- 1384 • Figure A.5 **(b)** shows that similar to the Psycho-B vs. ResNet-101 comparison in Figure 5 **(b)**, salience maps of PsychoDW’s low-mid level Phasor Blocks clearly emphasize specific object parts, while those of ConvNeXt-S are much more general and diffuse. Further examples of the former are shown in Figure A.6.
- 1385 • Figure A.7 shows that similar to for Psycho-B in Figure 8, PsychoDW’s SVC appears to  
1386 distribute object parts by scale between the three sub-bands, and individual filters within  
1387 each sub-band target distinct selections of object parts.

1388  
1389 Overall, these results are highly consistent with those for Psycho-B, showing that SVC abstractions  
1390 and object-part-centric Phasor Block representations also translate to ConvNeXt-S.



1418  
1419  
1420  
1421  
1422  
1423  
1424  
1425  
1426  
1427  
1428  
1429  
1430  
1431  
1432  
1433  
1434  
1435  
1436  
1437  
1438

Figure A.5: **(a)** Top spatial principal components of SVC filters learnt by PsychoDW trained on ImageNet-1K. Bilinear smoothing has been applied. **(b)** Comparison between activation maps (via KPCA-CAM) of PsychoDW and ConvNeXt-S for a range of layer depths. Real and imaginary components are denoted by  $\mathbb{R}$  and  $\mathbb{I}$ .



1439  
1440  
1441  
1442  
1443  
1444  
1445  
1446  
1447  
1448  
1449  
1450  
1451  
1452  
1453  
1454  
1455  
1456  
1457

Figure A.6: Assorted activation maps (via KPCA-CAM) for mid-level Phasor Blocks of PsychoDW. Real and imaginary components are denoted by  $\mathbb{R}$  and  $\mathbb{I}$ .

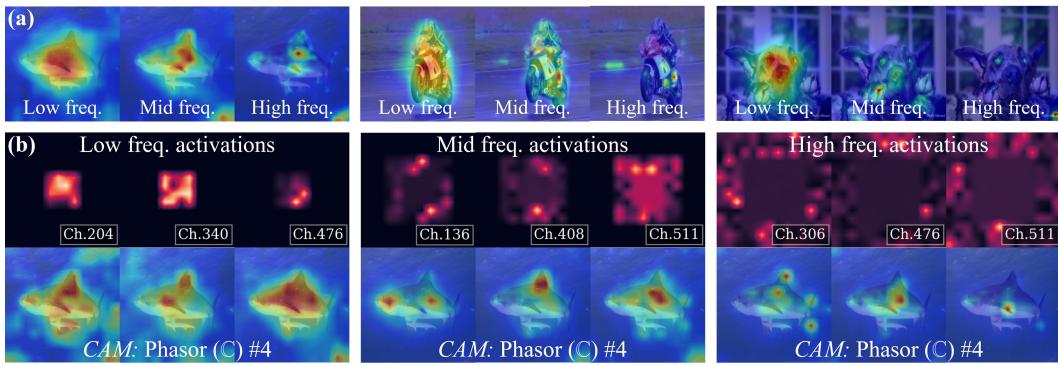


Figure A.7: PsychoDW Phasor Block salience maps (via HiResCAM) conditioned on gradients (a) from individual Spectral Branch sub-bands and (b) from individual frequency domain feature channels.

## C.2 SMALLER CLASSIFICATION DATASETS

Table A.10 presents experiment results for the CIFAR-10, CIFAR-100 and ImageNet-100 classification experiments.

Module	Parameters (M)	# Layers	CIFAR-10	CIFAR-100	ImageNet-100
ResNet50	25.56	54	94.14	78.10	80.90
Psycho-S	25.35	65	95.08	78.97	82.50
ResNet101	44.55	105	93.64	79.13	81.90
Psycho-B	42.01	95	94.99	79.49	83.60
ResNet152	60.10	156	93.17	77.51	83.60
Psycho-L	61.28	93	94.95	79.64	84.82
ResNet270	89.60	276	76.51	50.87	83.80
Psycho-H	88.61	93	94.68	79.89	85.00
ConvNeXt-S	50.22	113	94.09	76.96	86.98
PsychoDW	49.51	106	95.46	79.67	86.76

Table A.10: Classification results (% top-1 accuracies) for CIFAR-10, CIFAR-100 and ImageNet-100. Each pair of rows (separated by horizontal lines) compares a baseline CNN and the PsychoNet based on it.

**CIFAR-10** is a small scale dataset comprising 50000 natural images for training and 10000 images for testing across 10 classes, at a resolution of  $32 \times 32$  (Krizhevsky, 2009). For compatibility with this lower resolution (the ImageNet models have  $224 \times 224$  input resolution), we reduce initial down-sampling steps from our models. For ResNet and ResNet-based PsychoNet models, we removed the first maxpooling layer and set stride=1 for the first two ResBlocks that originally had stride=2. For ConvNeXt-S and PsychoDW, we replace the initial  $4 \times 4$  patch embedding layer with a standard  $3 \times 3$  Conv2D layer, and set stride=1 for the second downsampling layer. Table A.11 presents the training recipe for the CIFAR-10 experiments. Overall, all of our PsychoNet models outperformed their respective CNN baselines.

**CIFAR-100** contains the same images and train-test split as CIFAR-10, but with labels reorganised into 100 classes instead of 10. We use the same model configurations and training recipe as CIFAR-10, but increase the number of epochs to 90 since the greater number of classes results in a harder classification problem. Table A.11 presents the training recipe for the CIFAR-10 experiments. Overall, all of our PsychoNet models outperformed their respective CNN baselines.

**ImageNet-100** is a subset of the ImageNet dataset (Deng et al., 2009) that contains examples for 100 classes. It contains 130100 images for training and 5100 images for testing, at the original

1512  
1513  
1514 Table A.11: CIFAR-10 training recipe  
1515  
1516  
1517  
1518  
1519  
1520  
1521  
1522  
1523  
1524  
1525  
1526

Setting	Value
Image size	$32 \times 32$
Epochs	35
Batch size	64
Loss	Cross entropy
Optimizer	AdamW ( $\beta_1 = 0.9, \beta_2 = 0.999$ )
Scheduler	OneCycle
Learning rate (LR)	$10^{-3}$
Augmentation	crop, horizontal flip
GPU	1× NVIDIA A100: Psycho-S/B, ResNet50/101 1× NVIDIA H100: All other models

1527 resolution of  $224 \times 224$ . The model architectures remain the same as the ImageNet experiments,  
1528 but with the output linear layer modified to predict 100 logits. We use the same training recipe as  
1529 ImageNet-1K (Table A.8), but reduce the batch size to 128. Psycho-S/B and ResNet50/101 were  
1530 trained on 1× NVIDIA A100, while all other models used 1× AMD MI300X. Overall, the ResNet-  
1531 based PsychoNet models outperformed their respective baselines, but PsychoDW fell slightly short  
1532 of ConvNeXt-S.

1533  
1534 C.3 CLUSTERING VISUALISATION

1535 Figure A.8 presents an initial visualisation of clustering characteristics of Phasor Block activations  
1536 and SVC for Psycho-B. These show 2D PCA projections of features computed on samples from  
1537 10 randomly-selected classes from ImageNet-1K. Observable clustering emerges across both the  
1538 real and imaginary/magnitude-phase feature components, and becomes increasingly pronounced at  
1539 deeper layers.

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

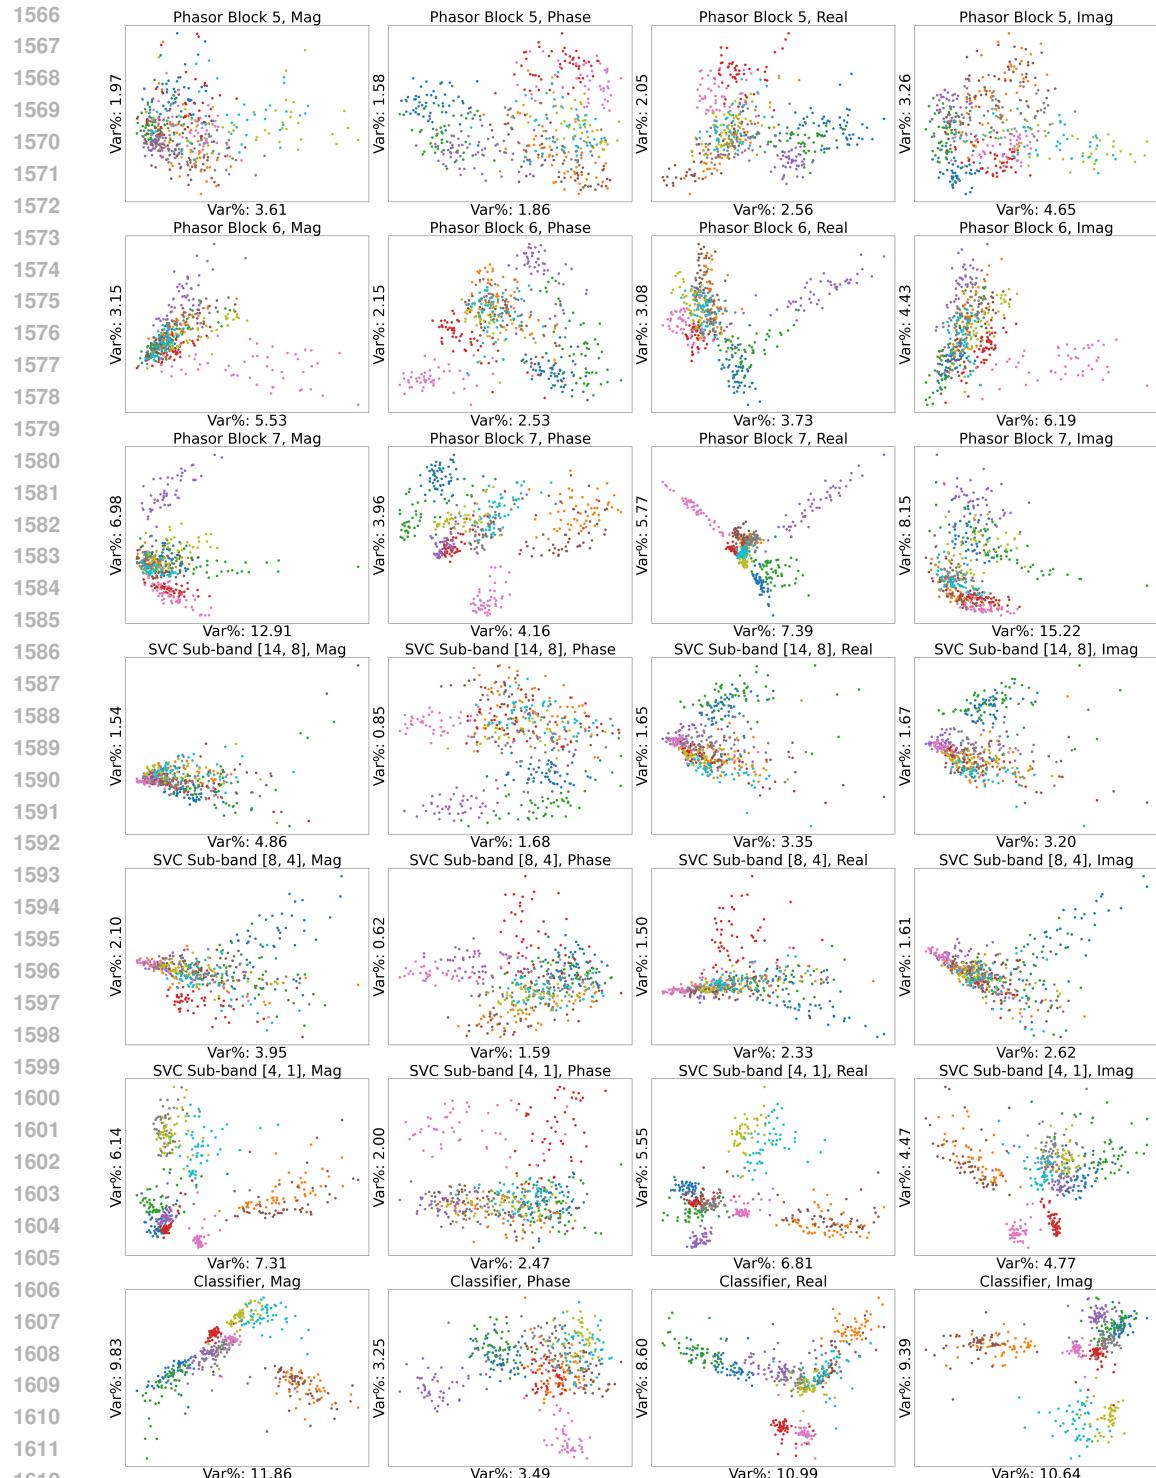


Figure A.8: Psycho-B clustering visualisation.

1620 **D ABLATION STUDIES**  
 1621

1622 **D.1 ARCHITECTURAL ABLATIONS**  
 1623

1624 We design four ablation model configurations to assess the impact of Phasor Blocks and SVC’s  
 1625 multiple Spectral Branches on classification performance and visual code quality.

1626 The postfix SP (Single Phasor) indicates that we remove all Phasor Blocks ( $\mathbb{C}$ )s and make up for  
 1627 the resultant layer and parameter deficit by adding additional ResBlocks. SP\_MB (Big/Large) were  
 1628 created by applying this modification to Psycho-B/L respectively, and in-depth architectural details  
 1629 of them are presented in Tables A.13 and A.14. The Single Branch (SB) models replace Spectral  
 1630 Branches with a single Hadamard Block with full band filters and no prior DropCrop operations.  
 1631 MP\_SB (Big/Large) and SP\_SB (Big/Large) were created by applying this modification to Psycho-  
 1632 B/L and SP\_MB (Big/Large) respectively. Table A.12 and Figure A.9 present quantitative and qual-  
 1633 itative results from this study.

1634 Spectral Branches appreciably improve classification accuracy (0.45-0.58%), except for between  
 1635 Big size SP\_MB and SP\_SB. In Figure A.9, we visualize the first two spectral bands of the MB  
 1636 models, as well as the corresponding bands isolated from the full-band filters of the SB models. The  
 1637 former are sparse and highlight distinct frequencies, while the latter exhibit a similar structure but  
 1638 with significant noise. This suggests that the explicit spectral decomposition of Spectral Branches is  
 1639 important for generating clear visual codes. Multiple Phasor Blocks slightly improve classification  
 1640 accuracy (0.24-0.252%) for MB models and have little effect on SB ones. However, Figure A.9  
 1641 shows that they drastically reduce noise and improve clarity of the SB filters, and moderately so for  
 1642 the MB ones.

1643 Finally, we also try removing the Phasor ( $\mathbb{I}$ ) block from the Big size SP\_MB model, yielding a  
 1644 model without any Phasor Blocks (‘no\_phasor’). This further reduces accuracy slightly, and results  
 1645 in the filters exhibiting conjugate symmetry as shown in Figure 6. We further visualise salience  
 1646 maps for the no\_phasor ablation in Figure A.10. Compared to ResNet, the no\_phasor model places  
 1647 greater emphasis on edge information across intermediate layers, likely reflecting the downstream  
 1648 SVC making explicit use of high-frequency features. In later layers, salience increasingly concen-  
 1649 trates on object parts, including the shark’s head and fins, consistent with our broader evidence that  
 1650 SVC encourages object-part-based abstraction (Section 4). However, when Phasor Blocks are in-  
 1651 cluded—as in the full Psycho-B model—object parts are isolated far more clearly in the salience  
 1652 maps. This suggests that while SVC promotes part-focused semantic representations, the complex  
 1653 representations introduced by Phasor Blocks are important for expressing these abstractions cleanly  
 1654 and with high specificity.

1655 **Table A.12: Ablation study results.** We compare all combinations of MB/SB and MP/SP model  
 1656 configurations, for Big and Large model sizes, using ImageNet top-1 accuracy.

1657 Model	1658 Multiple Phasor 1659 Blocks	1660 Multiple (Spectral) 1661 Branches	1662 Top-1 Acc. (%) (Psycho-B base)	1663 Top-1 Acc. (%) (Psycho-L base)
1660 MP_MB (Psycho-B/L)	✓	✓	78.846	79.848
1661 MP_SB	✓	✗	78.394	79.268
1662 SP_MB	✗	✓	78.600	79.596
1663 SP_SB	✗	✗	78.548	79.124
1664 no_phasor (SP_SB w/o Phasor ( $\mathbb{I}$ ))	✗	✗	78.44	

1666 Note that as per He et al. (2016), ResBlocks each comprises  $1 \times 1$ ,  $3 \times 3$  and  $1 \times 1$  kernel size Conv2D  
 1667 layers. In the below architecture tables, we denote their respective output channel sizes with  $d_{in}$ ,  
 1668  $d_{bot}$  and  $d_{out}$  respectively (‘bot’ is short for bottleneck, as these layers follow a channel bottleneck  
 1669 configuration). We also write ‘stride=2’ if a ResBlock performs  $2 \times$  spatial downsampling, since it  
 1670 is achieved by setting stride=2 in the  $3 \times 3$  Conv2D layer.

1671  
 1672  
 1673

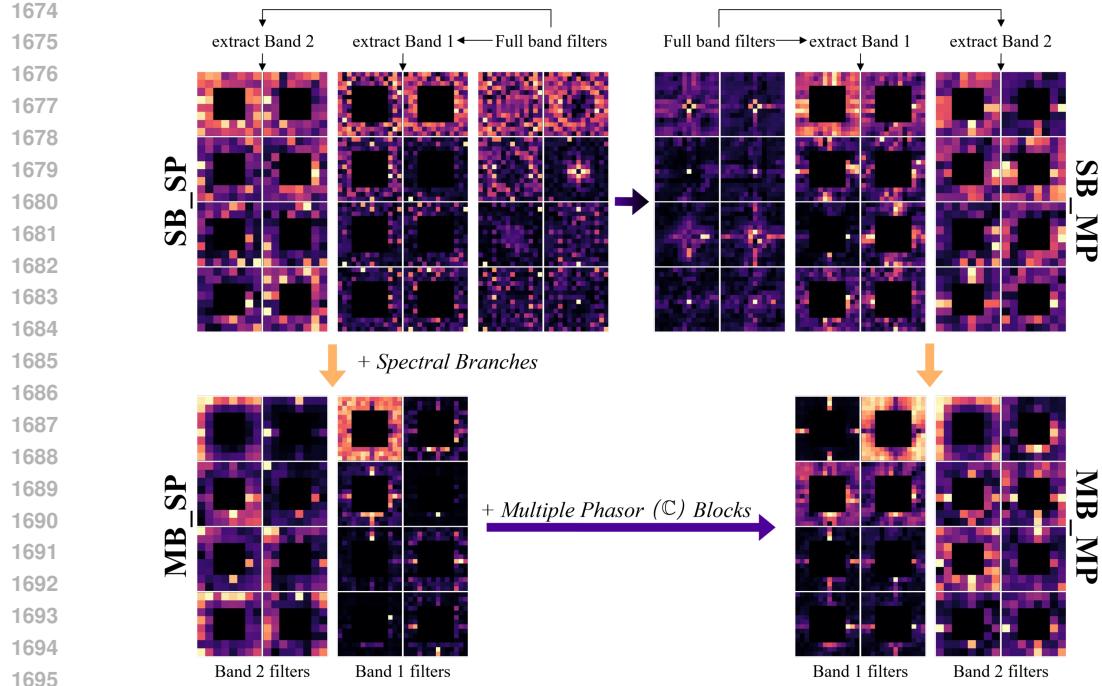


Figure A.9: Most significant channel-wise principal components of learnt spectral filters from Large size ablation models. We show the first two sub-bands of filters for MB models, and the full band filter for SB models.

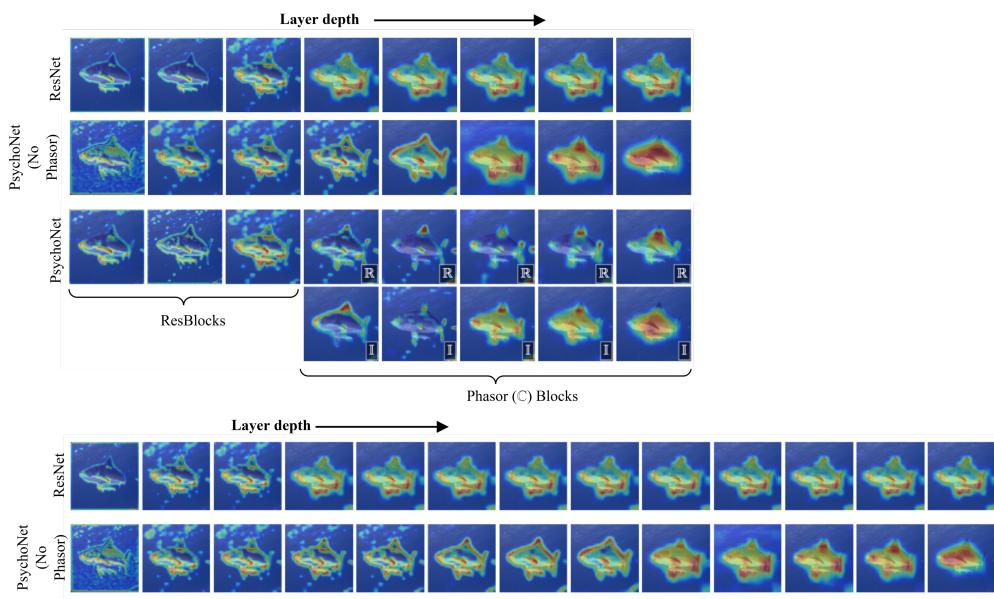


Figure A.10: (a) An extended version of Figure 5 (b) which additionally adds KPCACAM salience maps for the *no\_phasor* ablation model. (a) Additional salience maps comparison between ResNet-101 and *no\_phasor*.

### D.1.1 FLOP EFFICIENT MODELS

As we identified that PsychoNet used considerably more FLOPs than ResNet and ConvNeXt baselines (see Table A.12), we conducted the following ablation experiment to show a fairly straightfor-

1728 Table A.13: Detailed architecture of the SP\_MB (Big) ablation model.  
1729

1730 <b>SP_MB (Big) architecture</b>	
1731 Parameters (M)	42.263
1732 # Layers (overall)	91
1733 # Layers (complex)	2
1734 <b>Blocks</b>	
1735 Input layer	Conv2D( $7 \times 7$ , $d_{in}=3$ , $d_{out}=64$ , stride=2), MaxPool( $3 \times 3$ , stride=2)
1737	<b>56 × 56</b> : $[d_{in}=64, d_{bot}=256, d_{out}=256]$
1738	<b>56 × 56</b> : $[d_{in}=256, d_{bot}=64, d_{out}=256] \times 2$
1739	<b>28 × 28</b> : $[d_{in}=256, d_{bot}=128, d_{out}=512, \text{stride}=2]$
1740	<b>28 × 28</b> : $[d_{in}=512, d_{bot}=128, d_{out}=512] \times 7$
1741	<b>28 × 28</b> : $[d_{in}=512, d_{bot}=256, d_{out}=1024]$
1742	<b>28 × 28</b> : $[d_{in}=1024, d_{bot}=256, d_{out}=1024] \times 4$
1743	<b>14 × 14</b> : $[d_{in}=1024, d_{bot}=256, d_{out}=1024, \text{stride}=2]$
1744	<b>14 × 14</b> : $[d_{in}=1024, d_{bot}=384, d_{out}=1536]$
1745	<b>14 × 14</b> : $[d_{in}=1536, d_{bot}=384, d_{out}=1536] \times 6$
1746	<b>14 × 14</b> : $[d_{in}=1536, d_{bot}=128, d_{out}=512]$
1747	Phasor Blocks <b>14 × 14</b> : ( $\mathbb{I}$ ) $[d_{in}=512, d_{out}=512, \text{stride}=1]$
1748	SVC filters Sub-bands ([crop, drop]): $[14, 8], [8, 4], [4, 1]$ , $d_{filter} 512$
1749	Output layer Average pool, ComplexLinear( $d_{in}=1536, d_{out}=1000$ ), Softmax

1750 Table A.14: Detailed architecture of the SP\_MB (Large) ablation model.  
1751

1752 <b>SP_MB (Large) architecture</b>	
1753 Parameters (M)	60.42
1754 # Layers (overall)	90
1755 # Layers (complex)	2
1756 <b>Blocks</b>	
1757 Input layer	Conv2D( $7 \times 7$ , $d_{in}=3$ , $d_{out}=64$ , stride=2), MaxPool( $3 \times 3$ , stride=2)
1759	<b>56 × 56</b> : $[d_{in}=64, d_{bot}=256, d_{out}=256]$
1760	<b>56 × 56</b> : $[d_{in}=256, d_{bot}=64, d_{out}=256] \times 2$
1761	<b>28 × 28</b> : $[d_{in}=256, d_{bot}=128, d_{out}=512, \text{stride}=2]$
1762	<b>28 × 28</b> : $[d_{in}=512, d_{bot}=128, d_{out}=512] \times 5$
1763	<b>28 × 28</b> : $[d_{in}=512, d_{bot}=256, d_{out}=1024]$
1764	<b>28 × 28</b> : $[d_{in}=1024, d_{bot}=256, d_{out}=1024] \times 6$
1765	<b>14 × 14</b> : $[d_{in}=1024, d_{bot}=512, d_{out}=2048, \text{stride}=2]$
1766	<b>14 × 14</b> : $[d_{in}=2048, d_{bot}=512, d_{out}=2048] \times 7$
1767	<b>14 × 14</b> : $[d_{in}=2048, d_{bot}=128, d_{out}=512]$
1768	Phasor Blocks <b>14 × 14</b> : ( $\mathbb{I}$ ) $[d_{in}=512, d_{out}=512, \text{stride}=1]$
1769	SVC filters Sub-bands ([crop, drop]): $[14, 8], [8, 4], [4, 1]$ , $d_{filter} 512$
1770	Output layer Average pool, ComplexLinear( $d_{in}=1536, d_{out}=1000$ ), Softmax

1771  
1772 ward architectural modification can be applied to improve FLOP efficiency of PsychoNet. Models  
1773 Psycho-S-7×7 and Psycho-B-7×7 are variations of Psycho-S and Psycho-B which decrease the spa-  
1774 tial feature size of some Phasor Block layers, obtained via:

- 1775  
1776 • We add an additional downsampling stage by setting stride = 2 in one of the  $14 \times 14$  spatial  
1777 feature size Phasor Blocks, so that the subsequent layers in the model now have a spatial  
1778 feature size of  $7 \times 7$ .
- 1779  
1780 • To preserve  $14 \times 14$  SVC filters, input features to sub-bands covering features from  $7 \times 7$   
1781 to  $14 \times 14$  sub-band are taken from the output of the  $14 \times 14$  Phasor Block immediately  
preceding the new downsampling stage.

As shown in Table A.15, both Psycho-S-7×7 and Psycho-B-7×7 slightly underperform the original PsychoNet models, but match the much lower FLOP counts of the baseline ResNets and beat their performance. This suggests that exploiting higher-resolution features from earlier stages while reducing spatial resolution in deeper models is an effective and scalable strategy for improving efficiency, which will be investigated further in future work. Full architectural configurations for Psycho-S-7×7 and Psycho-B-7×7 are presented in Tables A.16 and A.17 respectively.

Model	Parameters (M)	# Layers	FLOPs (G)	ImageNet-100 Top-1 Acc (%)
ResNet50	23.71	54	8.18	80.24
Psycho-S	22.59	62	12.30	82.50
Psycho-S-7×7	24.80	62	8.17	82.34
ResNet101	42.71	105	15.60	82.58
Psycho-B	39.24	90	30.12	83.60
Psycho-B-7×7	38.92	90	12.97	83.50

Table A.15: Comparison of FLOP-efficient Psycho-7×7 models against PsychoNet and ResNet baselines on ImageNet-100 classification.

Table A.16: Detailed architecture of Psycho-S-7×7.

Psycho-S-7×7	
Parameters (M)	24.80
# Layers (overall)	62
# Layers (complex)	9
GFLOPs	8.17
Blocks	
Input layer	Conv2D( $7 \times 7$ , $d_{in}=3$ , $d_{out}=64$ , stride=2), MaxPool( $3 \times 3$ , stride=2)
Initial CNN layers	ResNet50 stem and first two resolution stages: 3× ResBlocks at $56 \times 56$ , 4× ResBlocks at $28 \times 28$
Phasor Blocks	<p><b>28 × 28:</b> (I) [<math>d_{in}=128</math>, <math>d_{out}=256</math>, stride=2]  <b>14 × 14:</b> (C) [<math>d_{in}=256</math>, <math>d_{out}=256</math>]  <b>14 × 14:</b> (C) [<math>d_{in}=256</math>, <math>d_{out}=384</math>]  <i>uparrow features used for 14 × 14 SVC input</i>  <b>7 × 7:</b> (C) [<math>d_{in}=384</math>, <math>d_{out}=512</math>, stride=2]  <b>7 × 7:</b> (C) [<math>d_{in}=512</math>, <math>d_{out}=512</math>]</p>
SVC filters	14 × 14 sub-bands ([crop, drop]): [14, 8], $d_{filter} = 384$ 7 × 7 sub-bands: [7, 4], [4, 1], $d_{filter} = 512$
Output layer	Average pool, ComplexLinear( $d_{in}=1408$ , $d_{out}=1000$ ), Softmax

1836  
 1837  
 1838  
 1839  
 1840  
 1841  
 1842  
 1843  
 1844  
 1845  
 1846  
 1847  
 1848  
 1849  
 1850  
 1851  
 1852

Table A.17: Detailed architecture of Psycho-B-7×7.

Psycho-B-7×7		
1855	Parameters (M) 38.92	
1856	# Layers (overall) 90	
1857	# Layers (complex) 13	
1858	GFLOPs 12.97	
1859	<b>Blocks</b>	
1860	Input layer Conv2D( $7 \times 7$ , $d_{in}=3$ , $d_{out}=64$ , stride=2), MaxPool( $3 \times 3$ , stride=2)	
1861	Initial CNN layers ResNet101 stem and first two resolution stages: $3 \times$ ResBlocks at $56 \times 56$ , $4 \times$ ResBlocks at $28 \times 28$	
1863	<b>28 × 28:</b> (I) [ $d_{in}=128$ , $d_{out}=256$ , stride=2]	
1864	<b>14 × 14:</b> (C) [ $d_{in}=256$ , $d_{out}=256$ ]	
1865	<b>14 × 14:</b> (C) [ $d_{in}=256$ , $d_{out}=256$ ]	
1866	<b>14 × 14:</b> (C) [ $d_{in}=256$ , $d_{out}=384$ ]	
1867	<b>14 × 14:</b> (C) [ $d_{in}=384$ , $d_{out}=384$ ]	
1868	<b>14 × 14:</b> (C) [ $d_{in}=384$ , $d_{out}=384$ ]	
1869	↑ features used for $14 \times 14$ SVC input,	
1870	<b>7 × 7:</b> (C) [ $d_{in}=384$ , $d_{out}=512$ ]	
1871	<b>7 × 7:</b> (C) [ $d_{in}=512$ , $d_{out}=512$ ]	
1872	<b>7 × 7:</b> (C) [ $d_{in}=512$ , $d_{out}=512$ ]	
1873	SVC filters $14 \times 14$ sub-bands ([crop, drop]): [14, 8], $d_{filter} = 384$	
1874	$7 \times 7$ sub-bands: [7, 4], [4, 1], $d_{filter} = 512$	
1875	Output layer Average pool, ComplexLinear( $d_{in}=1408$ , $d_{out}=1000$ ), Softmax	

1876  
 1877  
 1878  
 1879  
 1880  
 1881  
 1882  
 1883  
 1884  
 1885  
 1886  
 1887  
 1888  
 1889