



A High-Throughput Screening Approach to Discovering Good Forms of Biologically-Inspired Visual Representation

Nicolas Pinto¹, David Cox² and James DiCarlo¹

1. The McGovern Institute for Brain Research and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology
2. The Rowland Institute at Harvard



The Rowland Institute at Harvard
HARVARD UNIVERSITY

Take home message:

GPU Metaprogramming dramatically accelerates the discovery of bio-inspired vision models that beat state-of-the-art computer vision systems

Hardware	CPUs				GPUs			
Manufacturer	Intel	Intel	Intel	NVIDIA	Sony, IBM, Toshiba	NVIDIA	NVIDIA	
Model	Q9450	Q9450	Q9450	7900 GTX	PlayStation 3	8800 GTX	GTX 280	
# cores used	1	4	4	4x96	2+6	4x128	4x240	
Implementation	MATLAB	MATLAB	SSE2	Cg	Cell SDK	CUDA	CUDA	
Year	2008	2008	2008	2006	2007	2007	2008	
Performance / Cost								
Full System Cost (approx.)	\$1,500**	\$2,700**	\$1,000	\$3,000*	\$400	\$3,000*	\$3,000*	
Relative Speedup	1x	4x	80x	544x	222x	1544x	2712x	
Relative Perf. / \$	1x	2x	120x	272x	833x	772x	1356x	

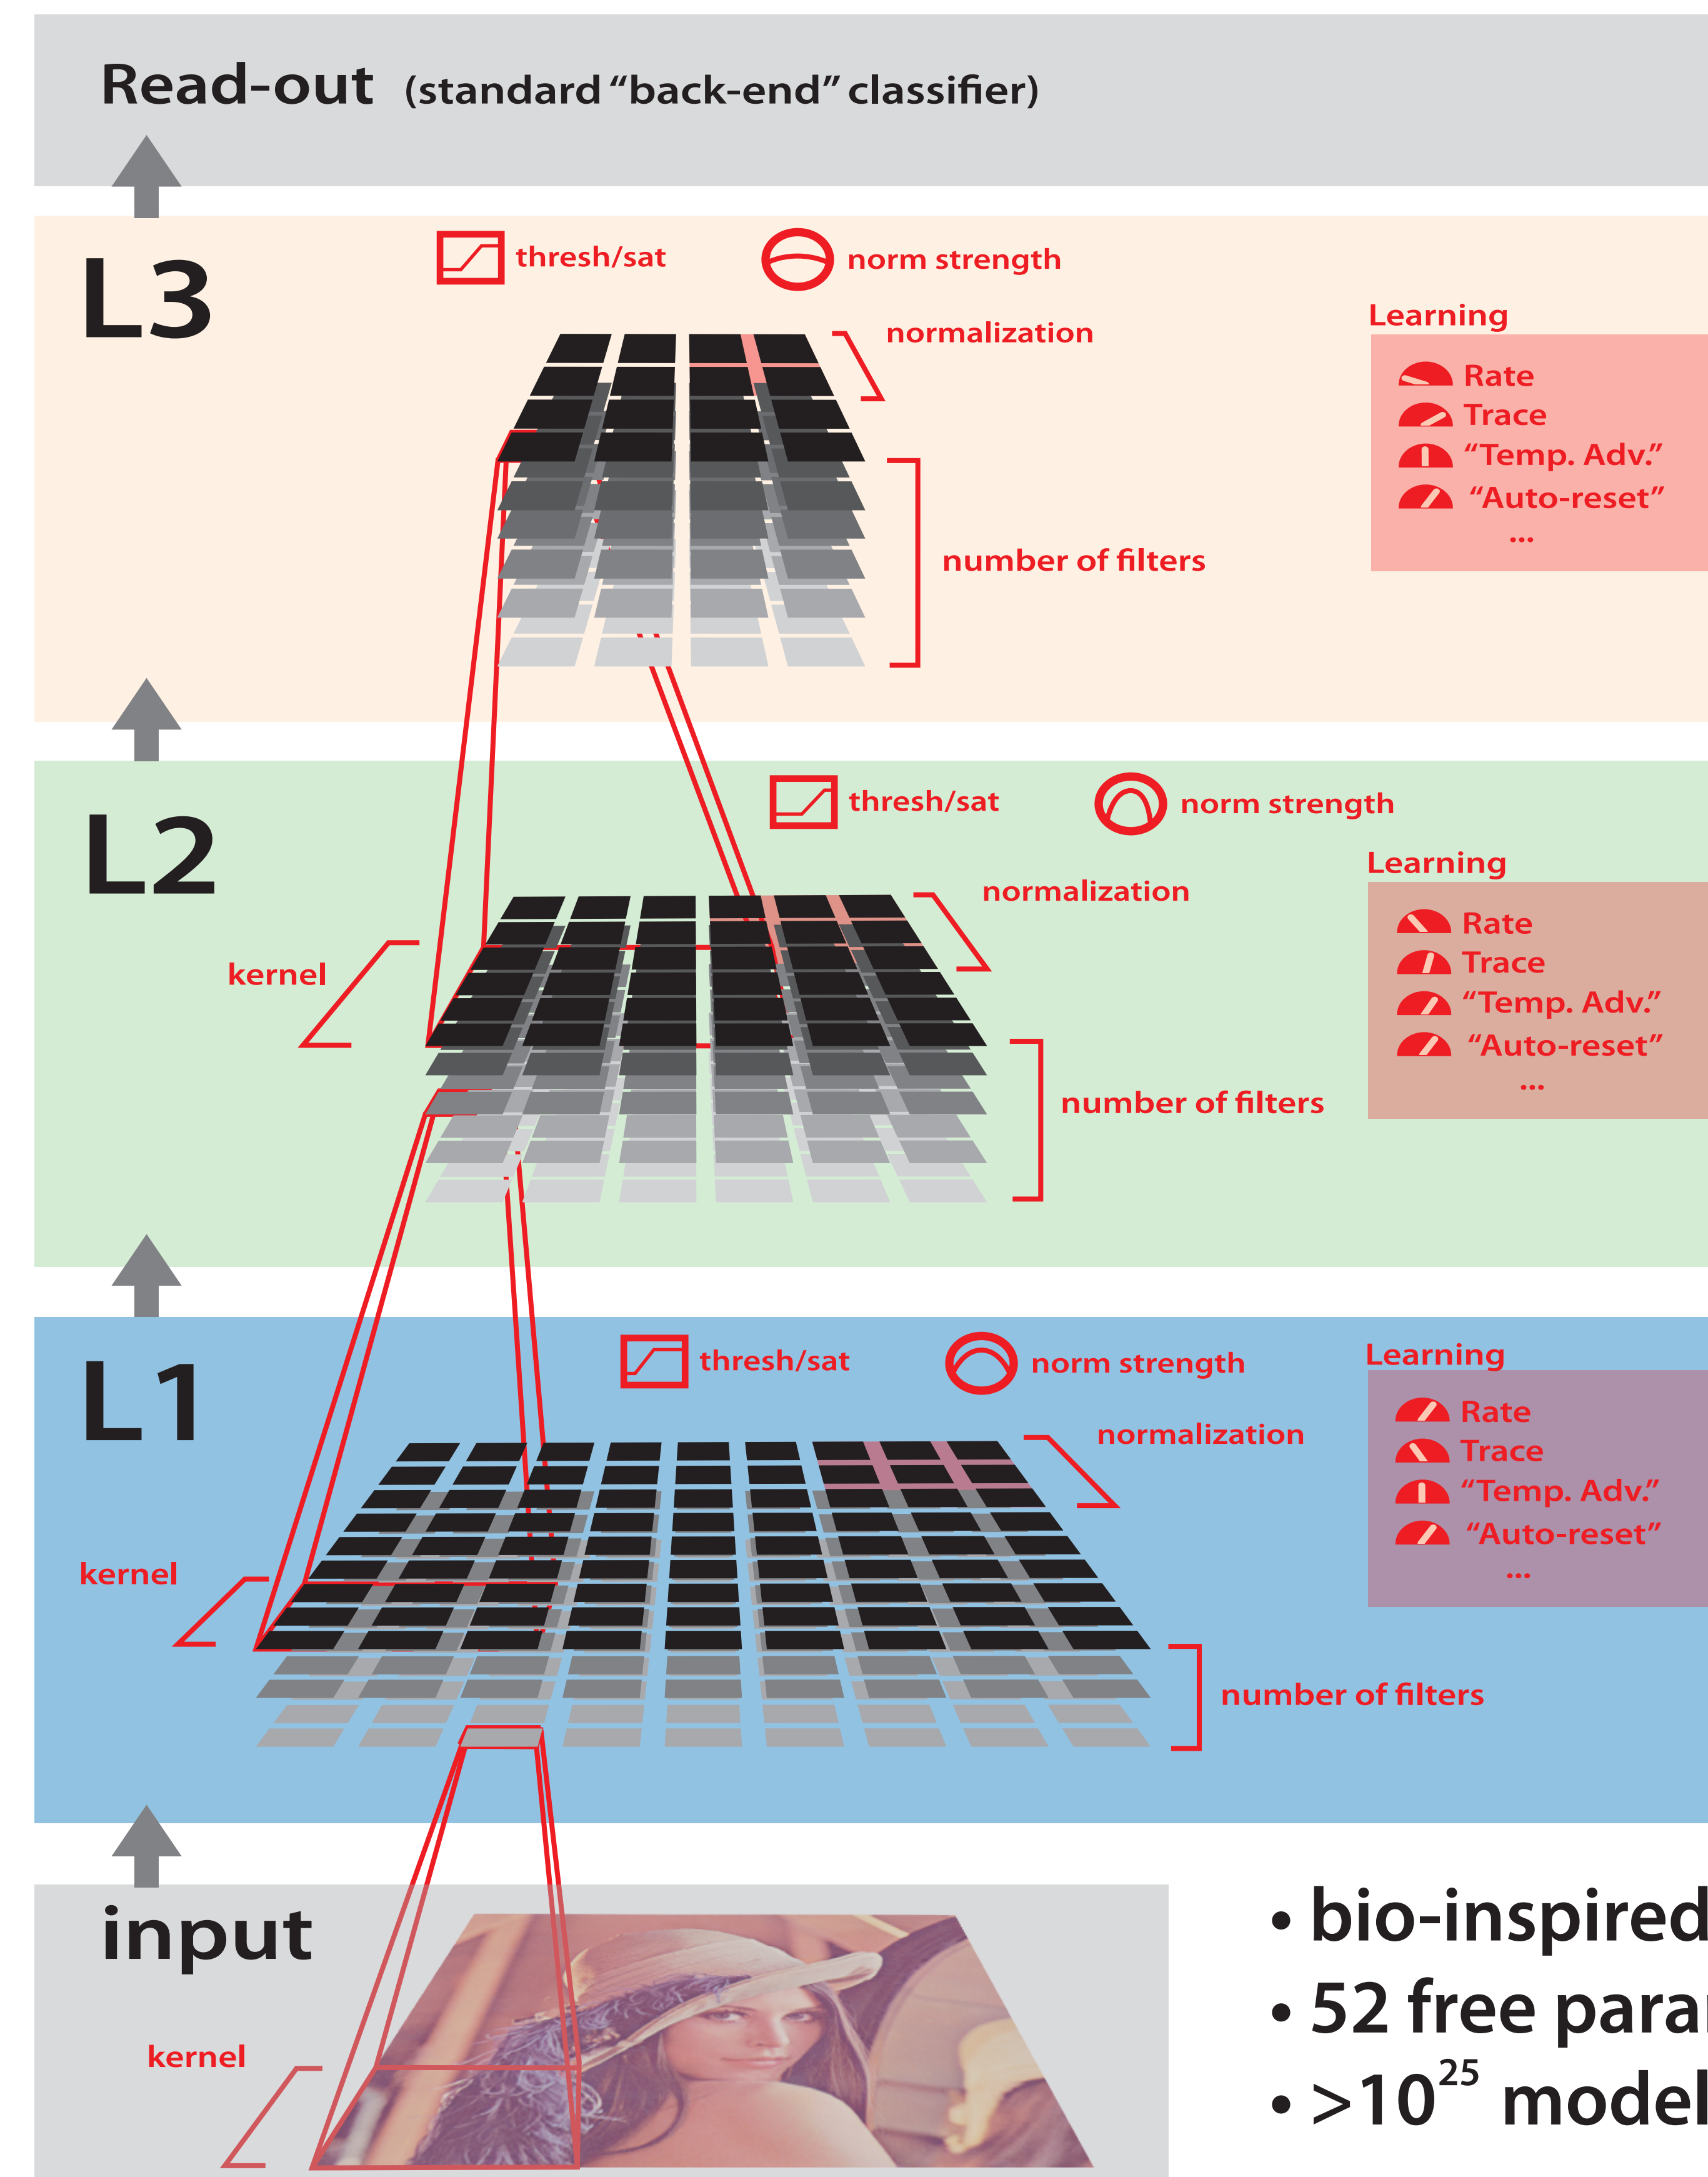
1 Abstract

While many models of biological object recognition share a common set of "broad-stroke" properties, the performance of any one model depends strongly on the choice of parameters in a particular instantiation of that model - e.g. the number of units per layer, the size of pooling kernels, exponents in normalization operations, etc. Since the number of such parameters (explicit or implicit) is typically large, and the computational cost of evaluating one particular parameter set is high, the space of possible model instantiations goes largely unexplored. Thus, when a model fails to approach the abilities of biological visual systems, we are left uncertain whether this failure is because we are missing a fundamental idea, or because the correct "parts" have not been tuned correctly, assembled at sufficient scale, or provided with enough training.

Here, we present a high-throughput approach to the exploration of such parameter sets, leveraging recent advances in stream processing hardware (high-end NVIDIA graphic cards and the PlayStation 3's IBM Cell Processor). In analogy to high-throughput screening approaches in molecular biology and genetics, we explored thousands of potential network architectures and parameter instantiations, screening those that show promising object recognition performance for further analysis. We show that this approach can yield significant, reproducible gains in performance in across an array of basic object recognition tasks, consistently outperforming a variety of state-of-the-art vision systems from the literature.

As the scale of available computational power continues to expand, we argue that this approach has the potential to greatly accelerate progress in both artificial vision and our understanding of the computational underpinning of biological vision.

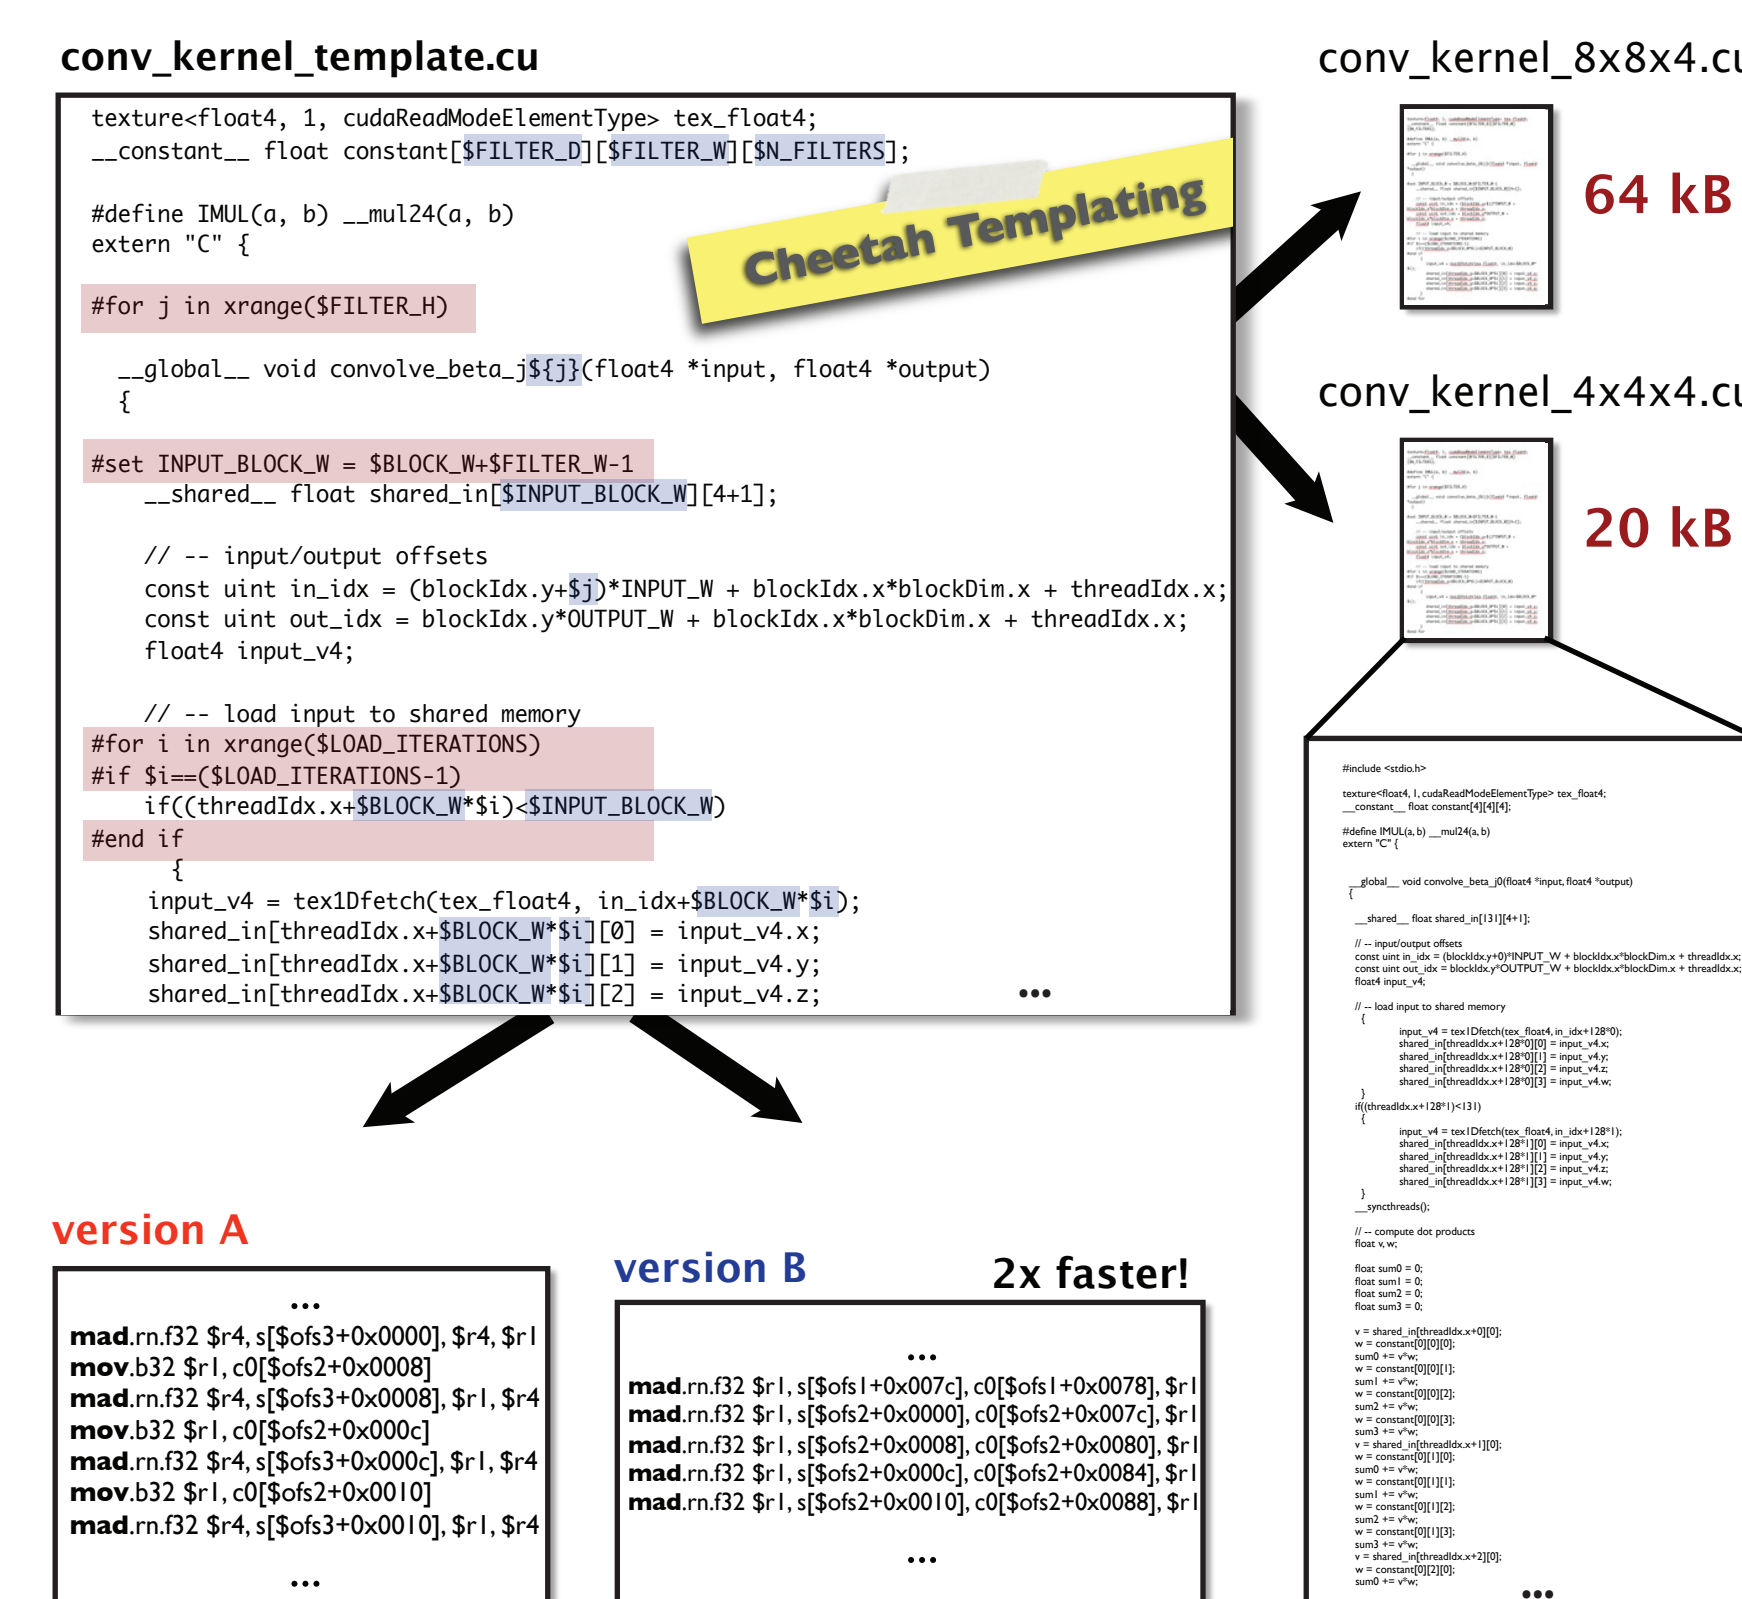
2 A Vast space of models to explore




- bio-inspired architecture
- 52 free parameters
- $>10^{25}$ models

3 Metaprogramming


- Recipe:**
- 1) Use scripting (e.g. Python) and templates,
 - 2) Instrumentalize your code,
 - 3) Let the computer auto-tune it!



4 GPU performance



5 High-throughput screening for good models




Unsupervised Learning

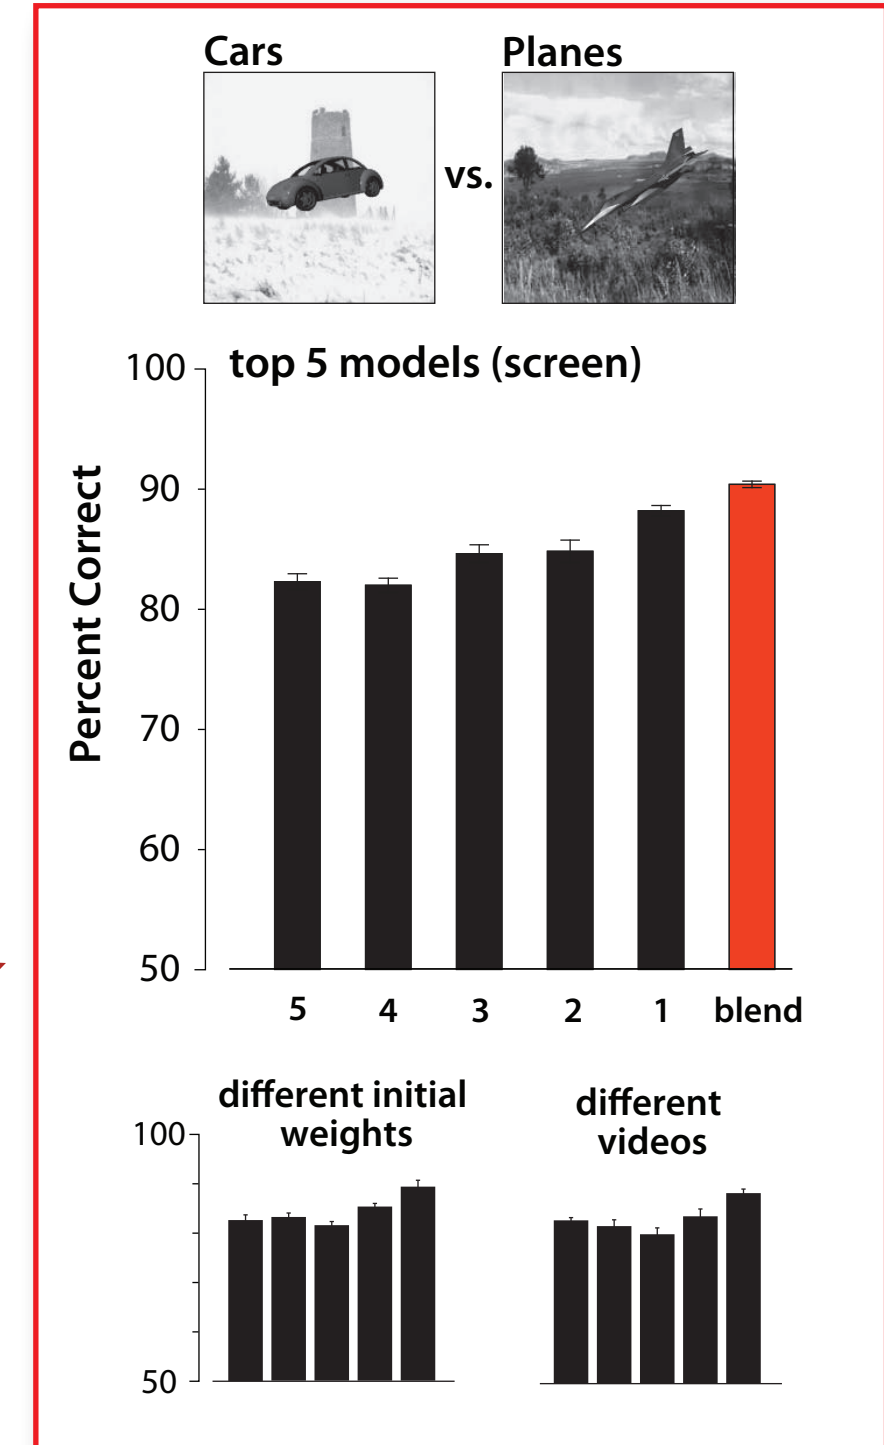


Law & Order "World"

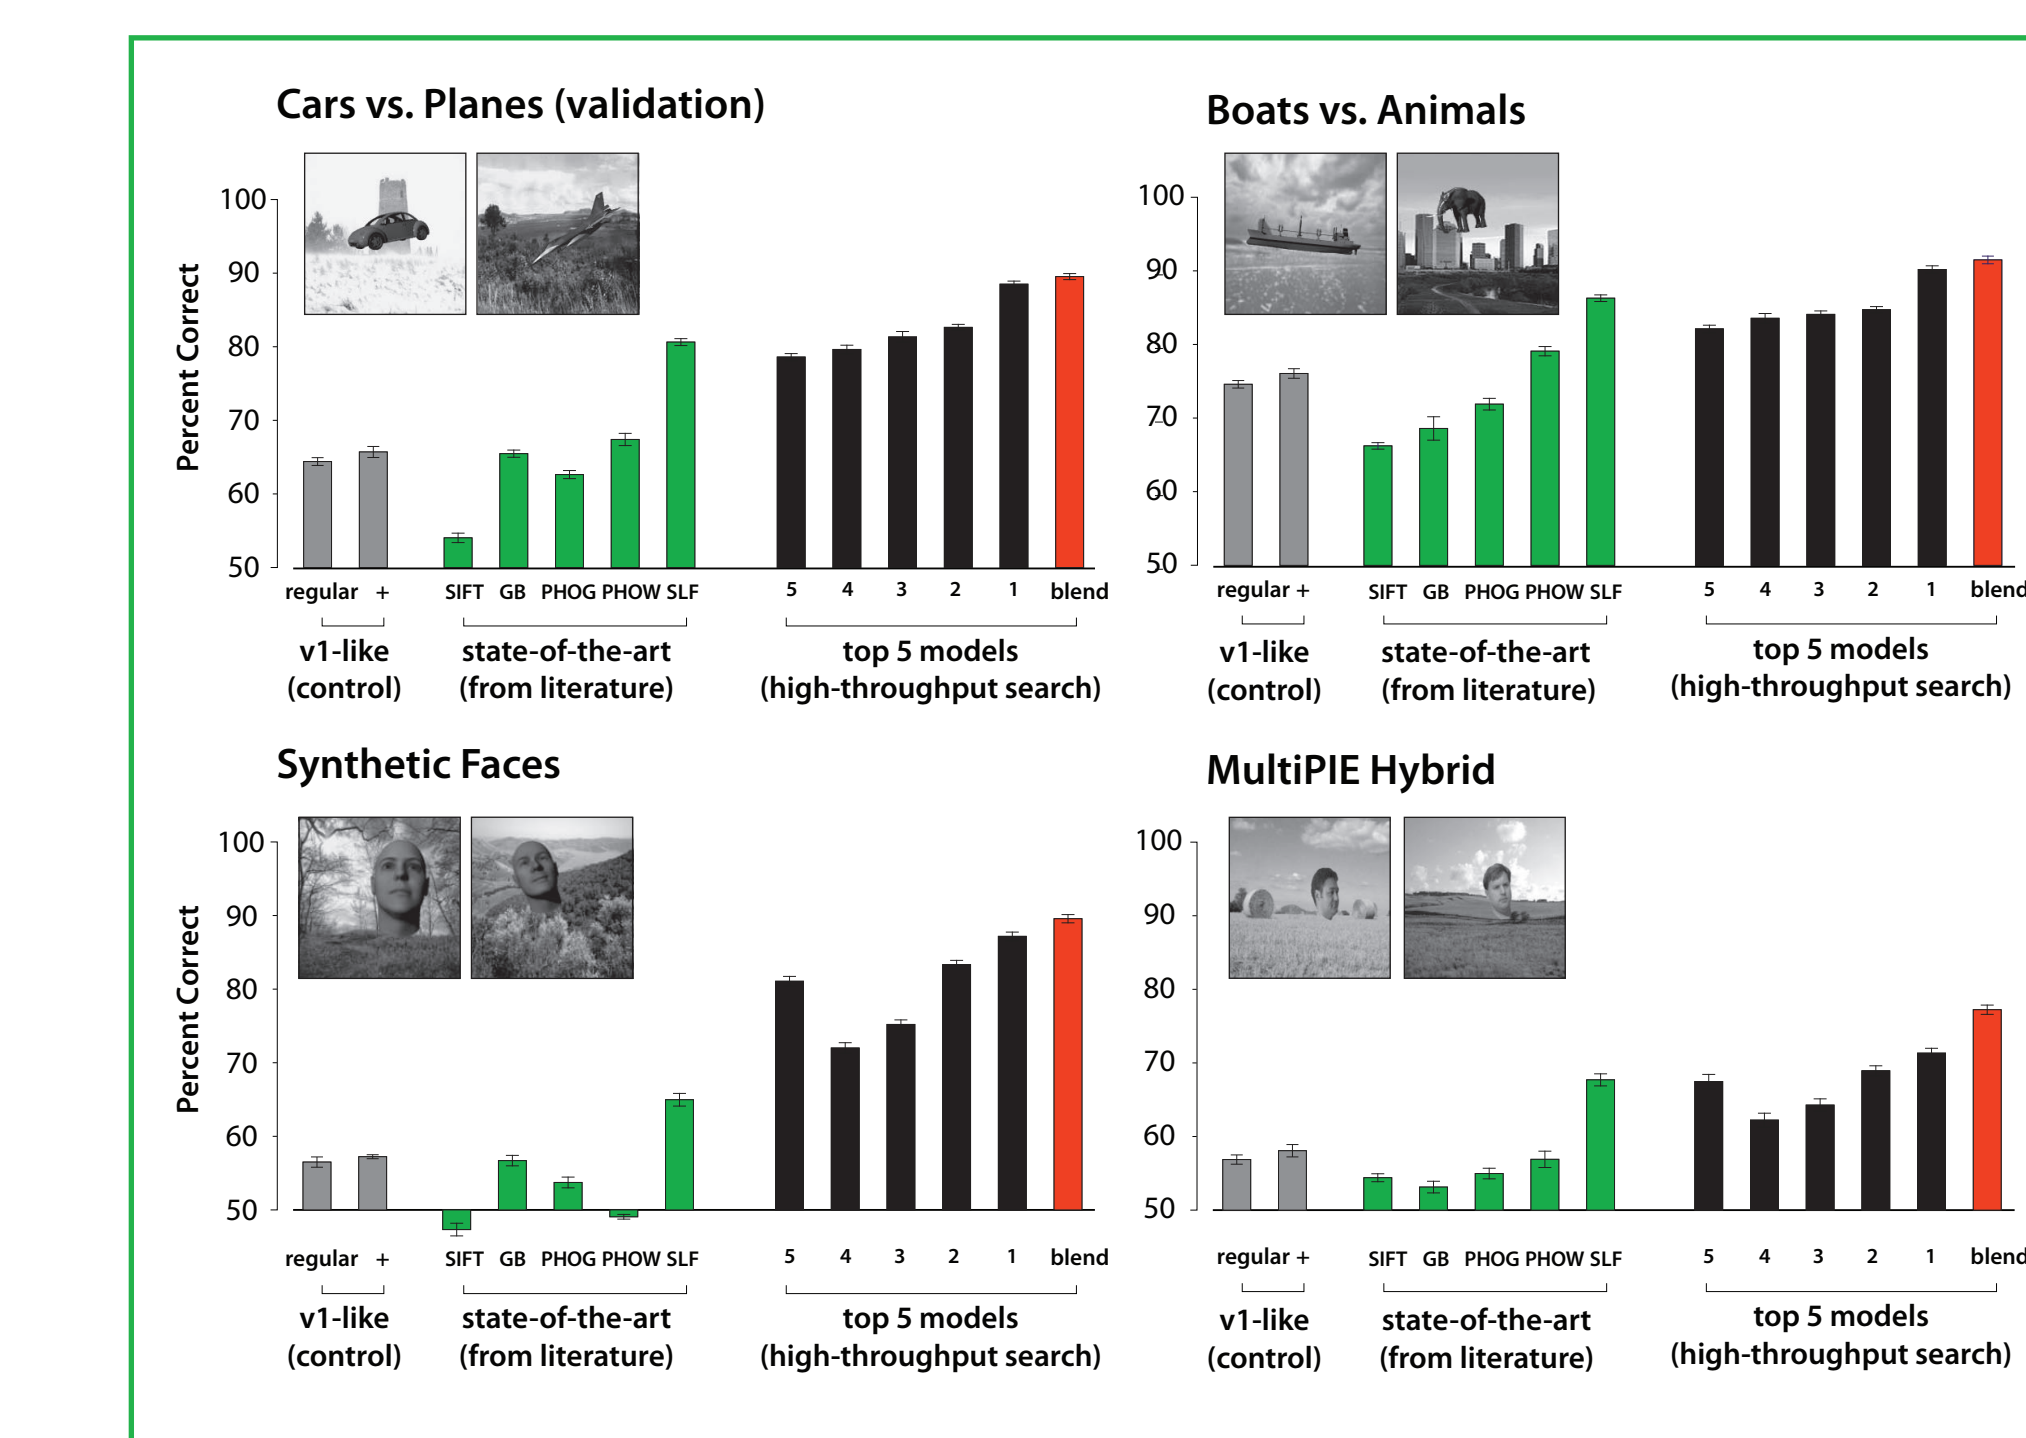
Screen Distribution



Best Models

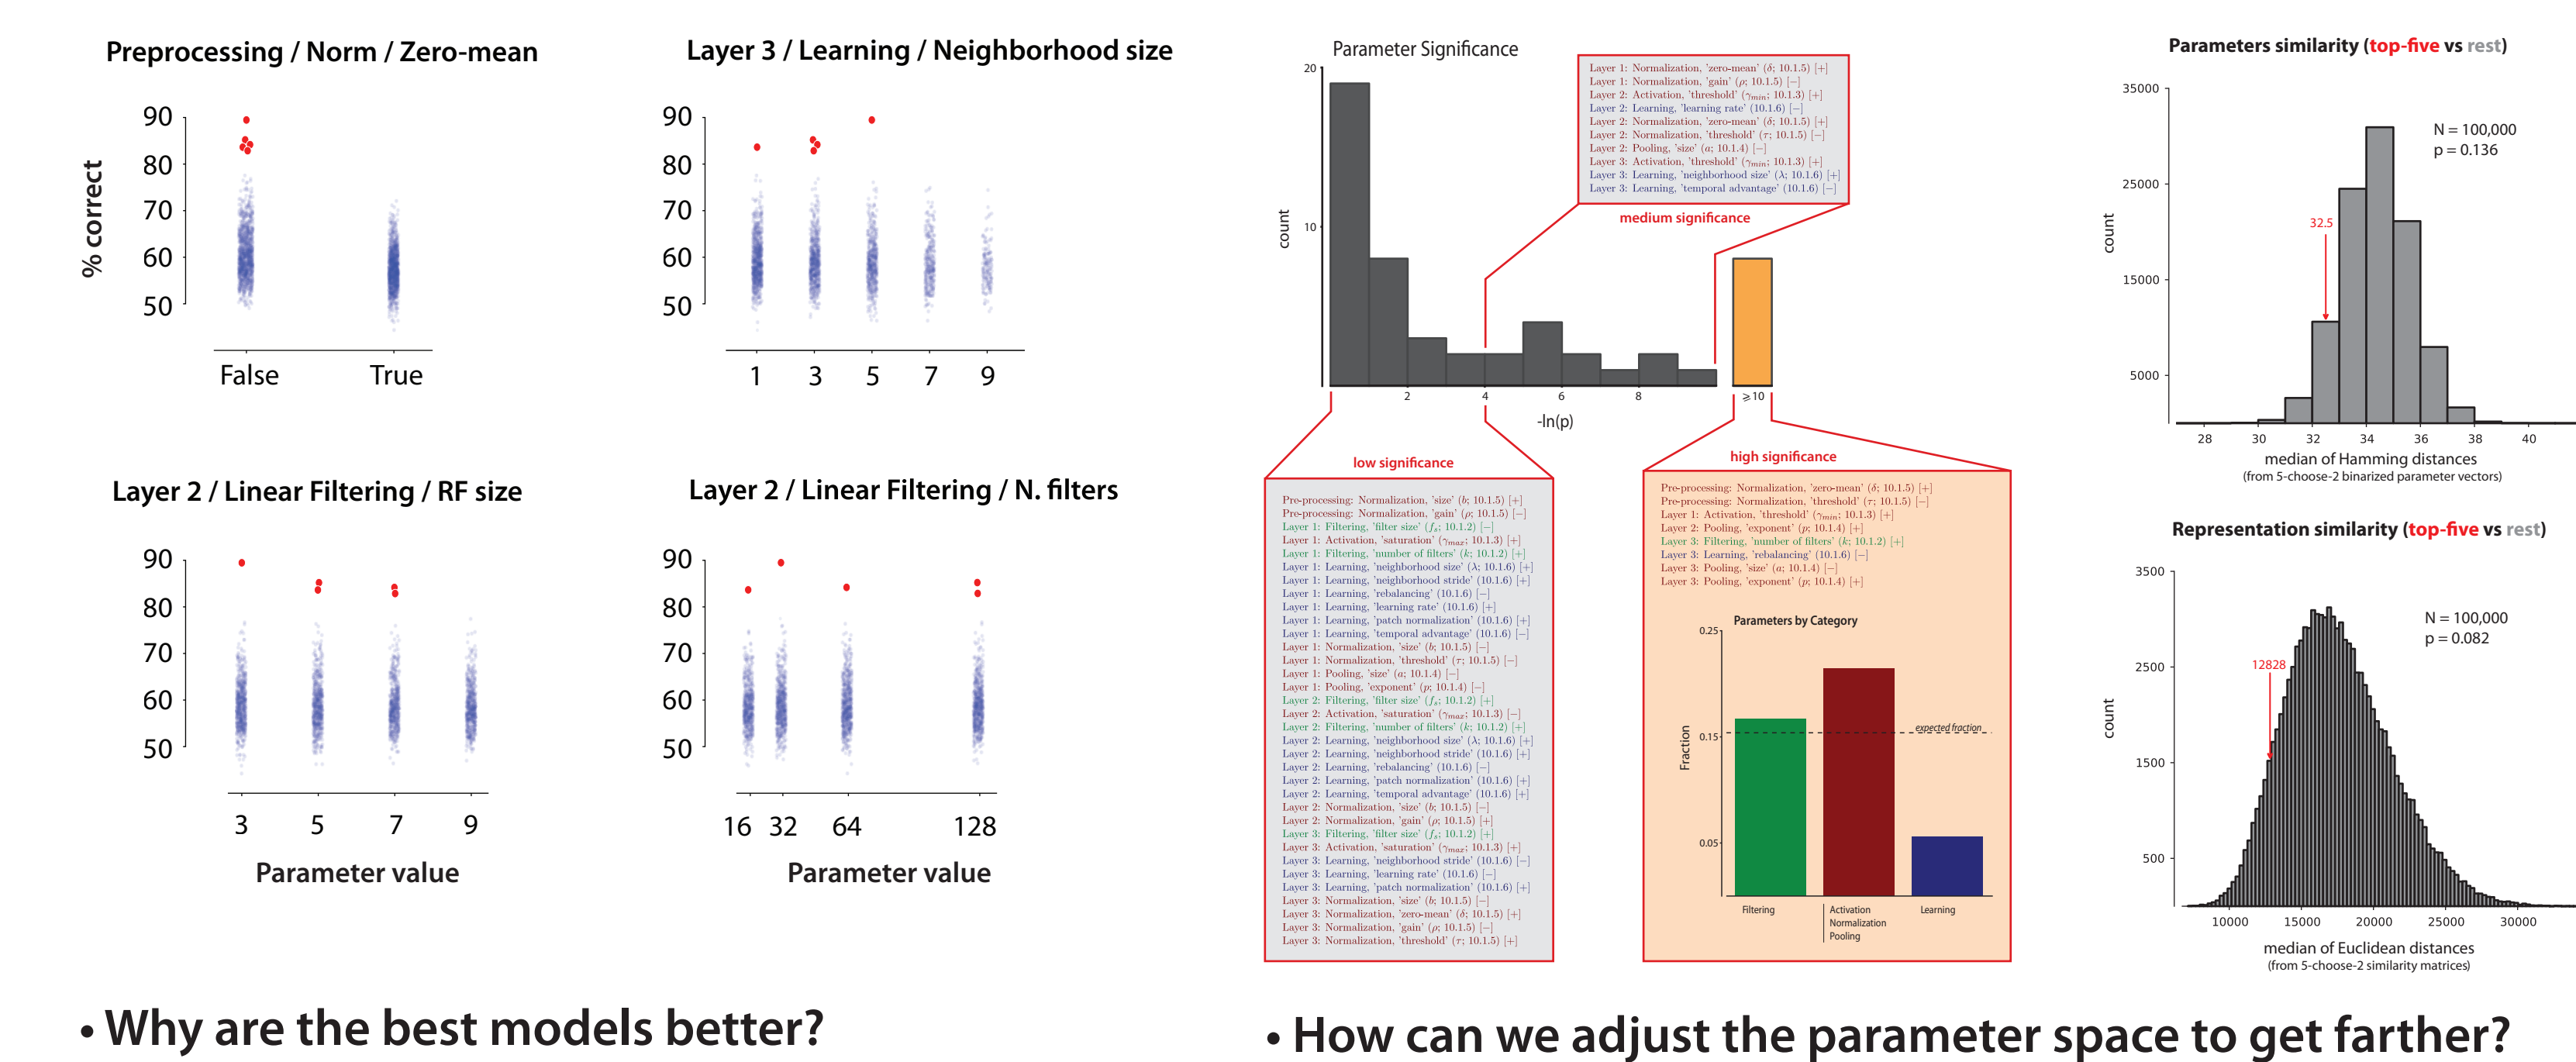


6 Result



We discovered models that beat state-of-the-art computer vision (object and face recognition)

7 Analysis: Why are these models good ?



• Why are the best models better?

• How can we adjust the parameter space to get farther?