

Article

A Deep Learning Model of Perception in Color-Letter Synesthesia

Joel R. Bock

Independent Researcher, 6736 Alamo Way, La Mesa, CA 91942, USA; sauerkraut@gmail.com

Received: 12 December 2017; Accepted: 8 March 2018; Published: 13 March 2018

Abstract: Synesthesia is a psychological phenomenon where sensory signals become mixed. Input to one sensory modality produces an experience in a second, unstimulated modality. In “grapheme-color synesthesia”, viewed letters and numbers evoke mental imagery of colors. The study of this condition has implications for increasing our understanding of brain architecture and function, language, memory and semantics, and the nature of consciousness. In this work, we propose a novel application of deep learning to model perception in grapheme-color synesthesia. Achromatic letter images, taken from database of handwritten characters, are used to train the model, and to induce computational synesthesia. Results show the model learns to accurately create a colored version of the inducing stimulus, according to a statistical distribution from experiments on a sample population of grapheme-color synesthetes. To the author’s knowledge, this work represents the first model that accurately produces spontaneous, creative mental imagery characteristic of the synesthetic perceptual experience. Experiments in cognitive science have contributed to our understanding of some of the observable behavioral effects of synesthesia, and previous models have outlined neural mechanisms that may account for these observations. A model of synesthesia that generates testable predictions on brain activity and behavior is needed to complement large scale data collection efforts in neuroscience, especially when articulating simple descriptions of cause (stimulus) and effect (behavior). The research and modeling approach reported here provides a framework that begins to address this need.

Keywords: synesthesia; deep learning network; color perception; generative adversarial network; cognitive modeling; character recognition; GPU computing

1. Introduction

1.1. Background

Synesthesia is a psychological phenomenon where sensory signals become mixed; input to one sensory modality produces an experience in a second, unstimulated modality [1]. For example, the experience of colors may be induced by seeing or hearing digits, letters or words. In “grapheme-color synesthesia”, viewed letters and numbers evoke mental imagery of colors. These color associations are involuntary, idiosyncratic and highly consistent over time [2]. The study of synesthesia has implications for furthering our understanding of brain architecture and function, as well as creative processes, language acquisition, and learning and memory performance [3]. The study of consciousness itself may benefit from investigation of synesthesia [4].

Due to the significance of synesthesia across disciplines, the scientific literature is rich with studies aiming to account for its physiological origins and behavioral manifestations. Investigations into cognitive and behavioral aspects and neurological substrates are reviewed by Hubbard and Ramachandran [1] and Rouw et al. [5].

Cognitive models have been proposed to identify partitions of cognition and perception in synesthesia. In [6], multiple interconnected pathways of form and color analysis under visual or

auditory stimuli are modeled. Symbol and color representation domains interact with semantic identification at higher stages of processing. An important contribution of the model in [6] is explicit description of the levels of inducer processing, from preconscious feature analysis to ultimate production of synesthetic assignment of color to the input.

Descriptions of the neurological basis of synesthesia center on two main theories. These theories posit that the synesthetic brain has either: (1) *cross-activation* between color processing area V4 and proximal visual word form area in the fusiform gyrus [7]; or (2) *disinhibited feedback* between circuits of bottom-up sensory input and higher-level visual areas [8]. Functional neuroimaging studies dynamically localize activated cortical regions in synesthetic perception versus controls [9], and provide valuable insight to validate and continue to develop such theories [10].

Rouw and co-workers [5] reviewed a number of studies elucidating the differences between synesthetes and controls based on functional MRI (fMRI) experiments. Their summary noted that six different regions of the brain were involved in overlapping research results, located in areas responsible for sensor, motor and attention and control processes. Synesthesia is clearly a very complex phenomenon, integrating neural activities and cognition associated with diverse functional regions of the brain.

1.2. Deep Learning Models

Deep learning facilitates machine learning from large scale data. Complicated, abstract representations of structure in computer vision, language processing, and many other domains can be explored [11,12]. Previously, computationally laborious problems in pattern recognition or artificial intelligence relied on the ability to engineer features from raw data that could be used to learn representation or mappings. Deep learning architectures develop internal representations naturally from this data, enabling new insights to emerge from empirical data for detection or classification applications [12].

Hinton demonstrated that multilayer generative models could learn the joint distribution of handwritten digit images and their labels [13]. These deep belief networks learned latent representations of the input in densely-connected hidden layers. Generative models were shown capable of: (1) learning low-level features in an unsupervised manner; and (2) learning very large numbers of parameters without over-fitting [13].

More recently, Goodfellow [14] introduced generative adversarial networks (GANs), a general framework for training deep learning networks. GANs eliminate the need for difficult probabilistic computations when learning hidden layer parameters.

The central idea of GANs is to establish competition between two deep network models—the *discriminator* (D) and the *generator* (G). G is tasked with generating samples $G(\mathbf{z})$ (drawn from $p_{\mathbf{z}}(\mathbf{z})$) that appear to D as having been drawn from the actual distribution $p_{data}(\mathbf{x})$. D must learn to discern between real data, and artificial data created by G . In this numerical game, model parameters are optimized alternatively to solve the minimax objective function

$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [\log(D(\mathbf{x}))] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))] \quad (1)$$

This formulation has an unique global optimum representing the real data distribution $p_{data}(\mathbf{x})$ [14]. This is true even when the prior distribution $p_{\mathbf{z}}(\mathbf{z})$ is random noise. Extensions of the GAN framework condition D and G on additional information such as class label [15,16].

1.3. Contribution of Present Study

A GAN is developed in the present work, in which both D and G are implemented as multilayer convolutional neural networks, as described in Section 2.3. The trained generator network G is a model for grapheme-color synesthesia.

Achromatic letter images, taken from database of handwritten characters [17], are used to stimulate the model and induce “computational” synesthesia. *G* learns to create a colored version of the inducing stimulus, according to a statistical distribution from experiments on grapheme-color synesthetes [18]. The identity of each symbol determines its concurrent color [8].

To the author’s knowledge, this work represents the first model that accurately produces spontaneous, creative mental imagery characteristic of the synesthetic perceptual experience. The GAN deep architecture can learn and (generalize) from gigabytes of example data in an efficient manner not feasible using alternative computing paradigms, without overfitting the training data. Several fundamental characteristics of color-letter synesthesia are present in this model. These include:

- Automatic induction of the psychological experience (in the trained GAN) [6];
- Invariance of colors across changes in form of inducing symbols [19];
- Stability of synesthetic associations over time [20]; and
- Learnability of synesthesia by non-synesthetes [21,22].

The current study is motivated to provide a template to reconcile the spectrum of experimental results on functional and structural aspects of this remarkable condition, as informed by behavioral and cognitive studies. Neuroimaging investigations collect huge amounts of data, even down to the granularity of single neuronal firings. These studies correlate regions and timings of brain activity under specific stimuli. Behavioral experiments begin with hypotheses and confirm, refute or modify these hypotheses after observation of results. The integration and interpretation of these two investigative fronts is a substantial objective.

Given the complexity of biological and cognitive processes involved in synesthesia, a path forward would benefit from a unifying computational approach. Significant scientific questions may be addressed moving forward following the suggestions of the current research. How do cognition and behavior arise from the interactions between activated neurons and assemblies under external sensory input? An ongoing challenge is to uncover causal relationships between “big neural” and “big behavioral” data [23]. Deep learning models are able to abstract neurological structures and representations at any desired level of detail, and process information given copious volumes of input data. The deep learning modeling and approach reported here provides a framework that begins to address this need.

2. Methods

2.1. Handwritten Letters Database

Handwritten letter images for training were extracted from the EMNIST dataset [17]. The raw images are stored as 28×28 pixels, in 8-bit integer format. A modeling sample comprising the EMNIST uppercase letters was constructed using the *By_Class* subset and annotations. We excluded lowercase letters and numeric digits resulting in 220,304 examples. Counts of individual letters in this sample varied from 2850 (letter ‘K’) to >29,000 (letter ‘O’); qualitatively, moderate variance in handwritten morphological structure for given letters was observed. No balancing of the 26 letter-classes was carried out; the aim here was not to develop a discriminative model for classification. (The EMNIST data can be obtained at: <https://www.nist.gov/itl/iad/image-group/emnist-dataset>.)

2.2. Synesthesia Color-Letter Pairs

Each grayscale letter image was converted to a 3-channel (*R,G,B*) image using experimental statistics of perceived colors in grapheme-color synesthesia as reported by Witthoft et al. [18]. The most frequently reported letter-color pairings from a large cohort of synesthetes ($n = 6588$) were used to represent the sample population (c.f. Figure 1 in [18]), recognizing that significant idiosyncratic differences in color experienced for a given letter exists between individuals. These aggregate “modal” colors for each letter were used to develop the basic examples for generative colorization model learning in the current study. These pairings are listed in Table 1.

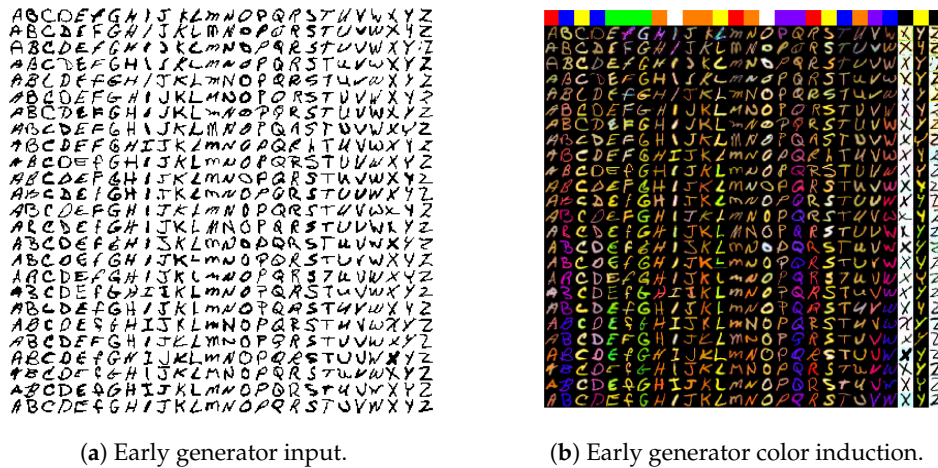


Figure 1. Synesthetic letter colorization by the evolving generative network: (a) generator input; and (b) generator output. Early training: 1st epoch; <3000 iterations. One letter per column. Iteration count progresses from top to bottom.

Table 1. Grapheme-color associations of 6588 synesthetes. The most common color assignment reported for each letter is shown. After [18].

Color	Letters
Red	A, M, R
Blue	B, D, T, W
Green	E, F, G
Yellow	C, L, S, Y
Orange	H, J, K, N, U
Purple	P, Q, V
White	I, O
Black	X, Z

2.3. Numerical Implementation

The conditional GAN model of grapheme-color synesthesia perception was adapted from a deep convolutional neural network (CNN) implementation described in [24]. The generator network encodes the input image by six successive hidden layers, each outputting a reduced-dimensional image relative to the preceding layer. The representation of features of the original input image is increasingly abstracted and noise-filtered after each encoding layer of processing [25].

For example, the encoding layers H decrease pixel resolution as follows: Input: 64×64 ($\times 3$ grayscale/color channels, for G , D respectively); $H1$: ($64 \times 64 \rightarrow 32 \times 32 \times 64$); $H2$: ($32 \times 32 \times 64 \rightarrow 16 \times 16 \times 128$); $H3$: ($16 \times 16 \times 128 \rightarrow 8 \times 8 \times 256$); $H4$: ($8 \times 8 \times 256 \rightarrow 4 \times 4 \times 512$); $H5$: ($4 \times 4 \times 512 \rightarrow 2 \times 2 \times 512$); $H6$: ($2 \times 2 \times 512 \rightarrow 1 \times 1 \times 512$).

Six decoding layers follow; these layers invert the compression operations by deconvolution producing the resulting image, which is 64×64 ($\times 3$ channels) in dimension. G outputs three channels of putatively colorized letters for assessment of accuracy by D .

The final output of the generator is an image with synthetic colorization. The discriminator has similar multilayer convolutions, ultimately outputting a reduced image and probabilities of the input image being real or synthetic. Additional details of the GAN architecture appear in [24]. (See [24] Appendix 1. The image colorization GAN adapted here is based on code from the repository: <https://github.com/sawhney-kartik/gan-colorization>).

We developed code to download, extract and pre-process EMNIST handwritten letters [17]; synthetically colorize the raw examples based on synesthesia statistics [18]; and to train, test, analyze

and visualize results of the generative learning process. Deep learning experiments were carried out using the TensorFlow software library and Python language API [26]. The models were developed on a CPU/GPU-based system (i5-3470 16 GB RAM; NVIDIA GeForce GTX 1050 Ti 4 GB on-board).

Transformations and processing for each training image followed the protocol of [24] (An additional pre-processing step taken here was to digitally upsample the EMNIST images from 28×28 to 64×64 pixel format): (1) convert example (R,G,B) image (described in Sections 2.1 and 2.2) to CIE Lab color space; (2) present the L (grayscale) channel to the input of the generative model G ; (3) output synthetic a,b channel data from G , and reconstruct a full three-channel (L,a,b) color image; (4) present real and artificial images to discriminator D ; (5) evaluate objective function (Equation (1)); and (6) backpropagate errors through both networks, updating weights using stochastic gradient descent.

Generative model training was performed for three epochs on a random sample of 125,000 images (three epochs, for a total 375,000 iterations) from the EMNIST dataset. Periodically, generated images were converted back to (R,G,B) color space and stored to disk for human observation and post-analysis.

3. Results

Results of synesthetic letter colorization by the generative network are presented in Figure 1. Figure 1a displays typical inputs and Figure 1b the corresponding outputs in the very early phases of model training (<3000 iterations).

Each panel contains samples for one handwritten letter per column. Iteration count progresses from top to bottom. The top-most row on the right-hand side shows the true modal color-letter pairings from the experimental distribution reported in [18] (Table 1). Letters are presented in random order and are unbalanced; the example count seen for each letter is therefore not uniform across the alphabet.

At this stage of learning, the generative network is not producing realistic colored images matching the true distribution. Several letters begin to appear to align with their synesthetic concurrents, but most others miss the mark. The letters often perceived as either black or white in color-grapheme synesthesia (I,O,X,Z) are already reproduced fairly well. This is intuitively correct, as the generator does not need to learn to produce color for 2/3 of the raw pixel data for achromatic exemplars.

Following additional optimization of the GAN, more realistic results are observed. Consider the images shown in Figure 2. After three epochs of training, the synesthesia model generates colored letters with high accuracy. The colors produced are nearly indistinguishable from the actual distribution. Accuracy of color reproduction in the 26×26 grid is over 99%; very few instances (G: row 7; F: row 18; K: row 2; O: row 15; Q: row 13; and Y: row 9) are incorrectly colored.

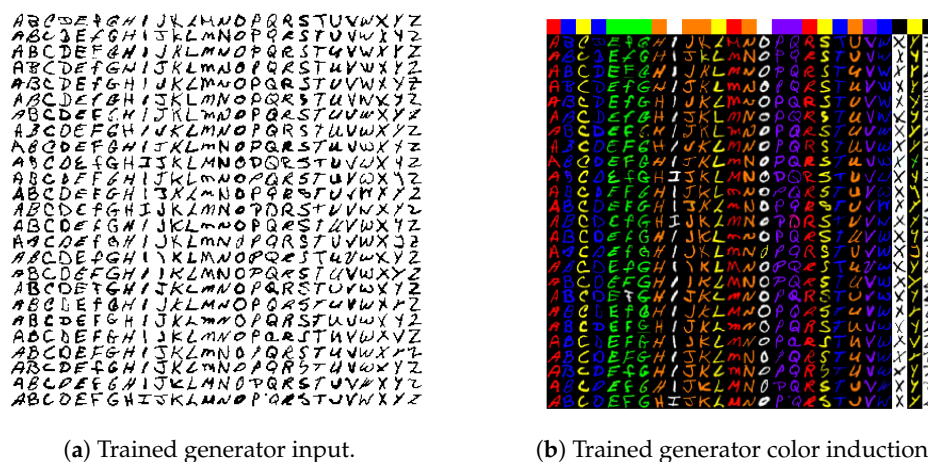


Figure 2. Synesthetic letter colorization by the trained generative network: (a) generator input; and (b) generator output. Late training: 3rd epoch, ~375,000 iterations. One letter per column. Iteration count progresses from top to bottom.

The results exhibited in Figure 2 are typical of those observed in out-of-sample tests on ~95,000 additional handwritten letter examples. Once the generator learns the true distribution of color–letter associations, the concurrent response remains consistent under additional stimuli.

4. Discussion

This research contributes the idea of applying a very general modeling paradigm (generative adversarial deep networks) to study synesthesia. We developed and applied a generative deep neural network to model perception in grapheme–color synesthesia. Grayscale letters are used to stimulate the model, which colorizes each letter. Training data were taken from experiments on a large sample of synesthetes [18], combined with a database of handwritten letters [17]. The spontaneous, creative mental imagery characteristic of the synesthetic perceptual experience is accurately reproduced by the model.

Automatic and consistent response to characters or digits are fundamental to synesthesia, even when the stimulus is non-physical (i.e., conceived in the mind’s eye) [19]. Recognition of the identity of a symbol determines its concurrent color [2,8]. In the present model, the generative network hidden layer weights encode information on pixel intensities describing the structure of each letter (and its identity) when perceived in two dimensions. We suggest that a similar mechanism may regulate the process of letter identification and association with color in synesthesia.

A model of synesthesia that generates testable predictions on brain activity and behavior is needed [2]. The experimental collection of “big data” in neuroscience leads to increased complexity of analysis and the distillation of conclusions, especially when articulating simple descriptions of cause (stimulus) and effect (behavior) [23]. The deep learning modeling and approach reported here represents a first step towards the integration of full scale fMRI functional neurological data with cognitive-behavioral experiments.

The present model reproduces the mechanism of character recognition and learned color association that characterizes color–letter synesthesia. After three presentations of each example image (in total, 375,000 iterations), highly accurate color–letter associations are generated by the model. We sought some confirmation of the number of letter presentations required to learn or induce synesthesia in the literature. In Colizoli et al. [21], non-synesthetes were trained by reading colored books, with four high-frequency letters presented in colored text. Subjects were subsequently shown to exhibit behavior indicative of synesthesia on a number of tests, including a modified “Stroop task” [27], a standard test for synesthetic effects. After reading a single artificially colored book (containing 49,000 words), the Stroop effect was observed (i.e., increased response times to letter–color paired stimuli incongruent with training pairs). Using this 49,000 word count threshold, the average number of letters per word in English (4.79) and the relative frequencies of occurrence of letters in the English language, a short simulation was programmed to estimate counts of each letter expected in a random sample of $49,000 \times 4.79$ letters, approximately one book in [21]. The simulation results suggest average expected training letter counts in [21] would be as follows:

[E = 29,792; T = 21,258; A = 19,014; O = 17,704; I = 16,391; N = 15,820; S = 14,904; H = 14,232; R = 13,959; D = 10,208; L = 9293; C = 6629; U = 6426; M = 5697; W = 5401; F = 5278; Y = 4762; G = 4604; P = 4585; B = 3492; V = 2318; K = 1852; J = 369; X = 326; Q = 217; Z = 179].

In the current deep learning model, after ≈ 3000 letter iterations, learning of synesthetic color–letter association was still incomplete. This is a lower threshold in the GAN model. It is not possible to draw definitive conclusions on statistical relevance to experiments in [21], but clearly something on the order of 10^4 colored examples (high frequency letters) are sufficient to learn synesthesia in behavioral experiments. The deep learning observations are roughly in accordance with these observations.

Three particularly ambitious areas in which to apply and extend this research are summarized below.

1. *Functional brain imaging.* Studies aimed to identify and localize structure and function within the synesthesia experience [5]. Experimental data from functional magnetic resonance imaging or other modalities could be used to build a complementary deep learning model with more explicit mapping of layers to physiologic modular components than used in the current study. Such a model could provide additional insights to refine the competing hypotheses of cross-activation (increased linkage between proximal regions) [7] or disinhibited feedback from higher-level cortical areas [8] in synesthesia. The architecture of the generative network used here comprised fully-connected, convolutional neural layers; different architectural designs could be developed (e.g., direct cross-wiring, or recurrent connections) to directly simulate and compare the relative accuracy and plausibility of these competing theories. This is beyond the current investigative scope.
2. *Language learning and memory.* Studying synesthesia may advance our understanding of human perception and information arrangement [6]. One theory proposed in [3] advances the idea that grapheme–color synesthesia develops in part by children to learn category structures; a fundamental task in literacy development is to recognize and discriminate between letters. Therefore, synesthesia might arise as an aid to memory. More generally, the ability to discern statistical regularities of printed letters or learn complex rules for letter combinations would assist learning at subsequent stages of literacy development [3]. In [28], small sample experiments suggested that additional sensory dimensions in synesthesia aid in memory tasks when compared to controls. We submit that the generative modeling approach of the current work may be useful to develop and test hypotheses in studies of language acquisition, memory and semantics.
3. *Consciousness studies.* Synesthesia can give insight into the neural correlates of consciousness, through interaction between sensory inputs and their mediation by semantics in the induction of phenomenal subjective experience [4]. Connecting neural activations with subjective aspects of consciousness (perception of shape, color, and movement of an object) is potentially achievable following a systematic experimental approach [29]. In deep learning, understanding representations within deep layers is easy at the layer level; at the level of individual neurons, such understanding is much more difficult [30]. Extensions of the deep learning model reported here may help to advance toward these formidable objectives.

Acknowledgments: The author thanks Tom Zingale, who provided the GPU-based computer used to conduct this research.

Conflicts of Interest: The author declares no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

GAN	Generative adversarial network
GPU	Graphics processing unit
CPU	Central processing unit
CNN	Convolutional neural network
fMRI	Functional magnetic resonance imaging

References

1. Hubbard, E.M.; Ramachandran, V. Neurocognitive Mechanisms of Synesthesia. *Neuron* **2005**, *48*, 509–520.
2. Mattingley, J. Attention, Automaticity, and Awareness in Synesthesia. *Ann. N. Y. Acad. Sci.* **2009**, *1156*, 141–167.
3. Watson, M.; Akins, K.; Spiker, C.; Crawford, L.; Enns, J. Synesthesia and learning: a critical review and novel theory. *Front. Hum. Neurosci.* **2014**, *8*, 98.
4. Van Leeuwen, T.M.; Singer, W.; Nikolić, D. The Merit of Synesthesia for Consciousness Research. *Front. Psychol.* **2015**, *6*, 1850.

5. Rouw, R.; Scholte, H.; Colizoli, O. Brain Areas Involved in Synaesthesia: A Review. *J. Neuropsychol.* **2011**, *5*, 214–242.
6. Rich, A.; Mattingley, J. Anomalous perception in synaesthesia: A cognitive neuroscience perspective. *Nat. Rev. Neurosci.* **2002**, *3*, 43–52.
7. Ramachandran, V.; Hubbard, E. Psychophysical investigations into the neural basis of synaesthesia. *Proc. R. Soc. Lond. B Biol. Sci.* **2001**, *268*, 979–983.
8. Grossenbacher, P.; Lovelace, C. Mechanisms of synesthesia: Cognitive and physiological constraints. *Trends Cogn. Sci.* **2001**, *5*, 36–41.
9. Jäncke, L.; Beeli, G.; Eulig, C.; Hänggi, J. The neuroanatomy of grapheme-color synesthesia. *Eur. J. Neurosci.* **2009**, *29*, 1287–1293.
10. Brang, D.; Hubbard, E.; Coulson, S.; Huang, M.; Ramachandran, V. Magnetoencephalography reveals early activation of V4 in grapheme-color synesthesia. *NeuroImage* **2010**, *53*, 268–274.
11. Bengio, Y. Learning deep architectures for AI. *Found. Trends Mach. Learn.* **2009**, *2*, 1–127.
12. LeCun, Y.; Bengio, Y.; Hinton, G. Deep Learning. *Nature* **2015**, *521*, 436–444.
13. Hinton, G.E.; Osindero, S.; Teh, Y.W. A Fast Learning Algorithm for Deep Belief Nets. *Neural Comput.* **2006**, *18*, 1527–1554.
14. Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems 27*; Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2014; pp. 2672–2680.
15. Mirza, M.; Osindero, S. Conditional Generative Adversarial Nets. *arXiv* **2014**, arXiv:1411.1784.
16. Gauthier, J. *Conditional Generative Adversarial Networks for Convolutional Face Generation*; Technical Report; Stanford University: Stanford, CA, USA, 2015.
17. Cohen, G.; Afshar, S.; Tapson, J.; van Schaik, A. EMNIST: An extension of MNIST to handwritten letters. *arXiv* **2017**, arXiv:1702.05373.
18. Witthoft, N.; Winawer, J.; Eagleman, D.M. Prevalence of Learned Grapheme-Color Pairings in a Large Online Sample of Synesthetes. *PLoS ONE* **2015**, *10*, e0118996.
19. Dixon, M.J.; Smilek, D.; Cudahy, C.; Merikle, P.M. Five plus two equals yellow. *Nature* **2000**, *406*, 365.
20. Cytowic, R. Synesthesia and mapping of subjective sensory dimensions. *Neurology* **1989**, *39*, 849–850.
21. Colizoli, O.; Murre, J.; Rouw, R. Pseudo-Synesthesia through Reading Books with Colored Letters. *PLoS ONE* **2012**, *7*, e39799.
22. Bor, D.; Rothen, N.; Schwartzman, D.; Clayton, S.; Seth, A. Adults Can Be Trained to Acquire Synesthetic Experiences. *Sci. Rep.* **2014**, *4*, 7089.
23. Sejnowski, T.; Churchland, P.; Anthony Movshon, J. Putting big data to good use in neuroscience. *Nat. Neurosci.* **2014**, *17*, 1440–1441.
24. Agrawal, M.; Sawhney, K. *Exploring Convolutional Neural Networks for Automatic Image Colorization*; Technical Report; Stanford University: Stanford, CA, USA, 2017.
25. Shwartz-Ziv, R.; Tishby, N. Opening the Black Box of Deep Neural Networks via Information. *arXiv* **2017**, arXiv:1703.00810.
26. Abadi, M.; Barham, P.; Chen, J.; Chen, Z.; Davis, A.; Dean, J.; Devin, M.; Ghemawat, S.; Irving, G.; Isard, M.; et al. TensorFlow: A system for large-scale machine learning. *arXiv* **2016**, arXiv:1605.08695.
27. Stroop, J.R. Studies of interference in serial verbal reactions. *J. Exp. Psychol.* **1935**, *18*, 643–662.
28. Mills, C.; Innis, J.; Westendorf, T.; Owsianiecki, L.; McDonald, A. Effect of a Synesthete's Photisms on Name Recall. *Cortex* **2006**, *42*, 155–163.
29. Crick, F.; Koch, C. A Framework for Consciousness. *Nat. Neurosci.* **2003**, *6*, 119–126.
30. Hinton, G. Learning multiple layers of representation. *Trends Cogn. Sci.* **2007**, *11*, 428–434.

