

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Investigations of neural networks and long-term memory

Permalink

<https://escholarship.org/uc/item/4zs311k0>

Author

Monk, Bradley

Publication Date

2021

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Investigations of neural networks and long-term memory

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Experimental Psychology and Cognitive Science

by

Bradley Ross Monk

Committee in Charge:

Professor Stephan Anagnostaras, Co-chair

Professor Roberto Malinow, Co-chair

Professor Christina Gremel

Professor Terry Sejnowski

Professor John Wixted

2021

Copyright

Bradley Ross Monk, 2021

All rights reserved.

The dissertation of Bradley Ross Monk is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2021

TABLE OF CONTENTS

Dissertation Approval Page	iii
Table of Contents	iv
List of Figures	v
List of Tables	vi
Acknowledgements	vii
Vita	viii
Abstract of the Dissertation	ix
General Introduction	1
Chapter 1: A Machine Learning Method to Identify Genetic Variants Potentially Associated with Alzheimer's Disease	8
Chapter 2: Synaptic stability and plasticity: a role of the actin filament network	25
General Discussion	65

LIST OF FIGURES

Figure 1.1: Q-Q plots of balanced ADSP dataset and Alzheimer's disease	11
Figure 1.2: Neural net prediction of case and control individuals	13
Figure 1.3: netSNP identifies potential AD-risk and AD-protective SNPs	15
Figure 1.4: netSNP validation	17
Figure 2.1: LTM models	29
Figure 2.2: Receptor density is a function of local diffusion rate	32
Figure 2.3: Metastable actin synaptic efficacy model	35
Figure 2.4: Simulated actin dynamics inside dendritic spines	38
Figure 2.5: Simulated actin dynamics and surface receptors	39
Figure 2.6: Dendritic spines are morphologically stable	42
Figure 2.7: Actin overexpression increases excitatory postsynaptic currents	45
Figure 2.8: Actin dissociates from thymosin during LTP	48

LIST OF TABLES

Table 1.1: netSNP Identified tSNPs with greatest absolute average CVt when APOE
locus variants were not included from the training set 12

ACKNOWLEDGMENTS

First and foremost, I would like to thank my advisors, Dr. Roberto Malinow and Dr. Stephan Anagnostaras for their invaluable mentorship throughout my graduate school career. Also my former advisor Jennifer Thomas, who was formative in preparing me for a doctoral training program. I would also like to thank my cohort for their friendship and camaraderie; in particular Scott Freeman who ventured with me on a startup project, which provided an opportunity for personal growth beyond our formal academic training. I'm forever indebted to my supportive lab mates, Chris, Kim, Stephanie, Sage, Mark, Chenyu, Jose, Karishma, Liz. I would like to acknowledge my lifelong friends Ben, Justin, Jamie, Chad, Tyler, my Beryl housemates Bry and Jan, my brother and side-kick Bryan and my darling sisters Vanessa and Desiree. I thank my partner Miren Edelstein for enriching my life immeasurably. Finally, I especially send thanks and love to my parents Jodi Monk and Ross Monk, words will never be enough to express my gratitude.

Chapter 1, in full, is a reprint of the material as it appears in *Frontiers in Genetics*. Monk, Bradley; Rajkovic, Andrei; Petrus, Semar; Rajkovic, Aleks; Gaasterland, Terry; Malinow, Roberto, 2021. The dissertation author was the primary investigator and author of this material.

Chapter 2, in part, is currently being prepared for submission for publication of the material. Monk, Bradley; Dore, Kim; Proulx, Christophe; Alphonso, Stephanie; Marino, Marc; Aronson, Sage; Malinow, Roberto. The dissertation author was the primary investigator and author of this material.

VITA

EDUCATION

Ph.D. in Experimental Psychology & Cognitive Science <i>University of California San Diego</i>	2021
M.A in Psychology <i>San Diego State University</i>	2011
B.S. Molecular Biology <i>San Diego State University</i>	2011
B.A. in Psychology (honors) <i>San Diego State University</i>	2006

PUBLICATIONS

- Monk, B.**, Rajkovic, A., Petrus, S., Rajkovic, A., Gaasterland, T., Malinow, R. (2021). Machine Learning Identifies Hundreds of Genes Linked to Alzheimer’s Disease. *Frontiers in Genetics*, 12.
- Shabel, S. J., Wang, C., **Monk, B.**, Aronson, S., & Malinow, R. (2019). Stress transforms lateral habenula reward responses into punishment signals. *Proceedings of the National Academy of Sciences of the United States of America*, 116(25), 12488–12493.
- Proulx, C. D., Aronson, S., Milivojevic, D., Molina, C., Loi, A., **Monk, B.**, Shabel, S. J., & Malinow, R. (2018). A neural pathway controlling motivation to exert effort. *Proceedings of the National Academy of Sciences of the United States of America*, 115(22), 5792–5797.
- Howell, K. K., **Monk, B. R.**, Carmack, S. A., Mrowczynski, O. D., Clark, R. E., & Anagnostaras, S. G. (2014). Inhibition of PKC disrupts addiction-related memory. *Frontiers in behavioral neuroscience*, 8, 70.
- Monk, B. R.**, Leslie, F. M., & Thomas, J. D. (2012). The effects of perinatal choline supplementation on hippocampal cholinergic development in rats exposed to alcohol during the brain growth spurt. *Hippocampus*, 22(8), 1750–1757.
- Thomas, J. D., Idrus, N. M., **Monk, B. R.**, & Dominguez, H. D. (2010). Prenatal choline supplementation mitigates behavioral alterations associated with prenatal alcohol exposure in rats. *Birth defects research. Part A, Clinical and molecular teratology*, 88(10), 827–837.

ABSTRACT OF THE DISSERTATION

Investigations of neural networks and long-term memory

by

Bradley Ross Monk

Doctor of Philosophy in Experimental Psychology & Cognitive Science

University of California San Diego, 2021

Professor Stephan Anagnostaras, Co-chair

Professor Roberto Malinow, Co-chair

This dissertation presents findings from two studies investigating the neuroscience of memory. Chapter 1 reports findings from a computational genomics study on Alzheimer's Disease (AD), a common neurodegenerative disease that causes memory loss and dementia. There is hope that genomic information can reveal insights to AD pathophysiology, along with aiding in risk assessment, screening, and diagnosis. This study involved using exome sequencing data from approximately 10,000 individuals (~5000 AD patients) to train neural net (NN) classifiers tasked with estimating the impact of single

genomic variants on AD polygenic risk. Together, these chapters and the studies presented therein add to our corpus of knowledge on memory systems, and contribute to our quest of understanding and treating dementia. Chapter 2 covers a study that details a model (supported by experimental results) that can explain how memories can be formed and maintained within neural network synapses, despite continuous and complete molecular turnover of synaptic proteins. Memory formation is thought to involve acute changes to synaptic weights; the ability for synaptic weights to remain stable over long time-periods and undergo evoked change is considered fundamental to our brain's information storage schema. Yet basic questions regarding synaptic plasticity remain unresolved, including a) how synaptic weights remain stable in the face of perpetual and complete molecular turnover; and b) how such weights can be modified to new stable levels by transient signals. Through a series of computational and biological experiments, we elucidate actin, a protein with unique polymer properties, as a potential central mediator of synaptic weights.

GENERAL INTRODUCTION

A topic that has captivated researchers in the broad fields of psychology and neuroscience, perhaps more than any other, is that of *memory*. The study of memory is vast and diverse, and includes for example research on the formation of place-preference memory (Gremel & Cunningham, 2010), fear memory (Anagnostaras, et al. 1999), recognition memory (Wixted, 2007), memory correlates in biological neural networks (Malinow, 2003) or artificial neural networks (Sejnowski, 2020). One reason why memory receives such widespread attention is that memory is a conserved ability that spans from humans down to single celled organisms. For example the actin cytoskeleton of single-celled amoebae exhibit bistable memory properties, allowing cells to orient and move towards chemoattractant signals (Pershin et al. 2009; Westendorf, et al. 2013; Artemenko et al. 2014; Skoge, et al., 2014). Individual cells in multicellular organisms have similar abilities. For example, white blood cells known as neutrocytes display long-lived directional memories, which are also attributed to cytoskeletal dynamics (Albrecht & Petty 1998; Prentice-Mott et al. 2016). Indeed, long-lived memory traces at the cellular level tend to follow a generic law centrally involving a dynamic network of filament scaffolding (Maiuri et al. 2015).

Structural plasticity also plays a role supporting higher-order forms of memory (e.g. explicit and implicit memory) stored within neural networks (Fischer et al. 1998; Okamoto et al. 2004). Information stored in neural networks is thought to be encoded by synaptic weights, which refers to the efficacy by which an upstream neuron evokes responses (often via postsynaptic glutamate receptors) in a downstream neuron at a particular synaptic connection (Kessels & Malinow, 2009; Huganir & Nicoll 2013). Converging evidence suggests structural plasticity at dendritic spines, femtoliter-sized sites of synaptic input, plays a central role in the induction and maintenance of synaptic weights (Halpain, 1998; Fukazawa, 2003; Kopec et al., 2006; Bosch, 2014). Compared to small spines, large spines generally have larger synapses (Bourne & Harris, 2012), more filamentous actin (Lin et al. 2005), more glutamate

receptors (Fischer et al. 1998), and greater synaptic weights (Malenka & Nicoll, 1999; Makino & Malinow, 2009; Nicoll & Roche, 2013). Also large spines are more robust, known to outlive small spines by at least hundreds of hours (Holtmaat, 2005).

While synaptic weights are considered fundamental to brain information storage, several important questions concerning these processes remain unresolved, including a) how synaptic weights remain stable despite protein turnover, and b) how such weights can be modified to new stable levels by transient signals. Chapter 2 of this dissertation addressed these questions by examining structural plasticity with a focus on actin dynamics in dendritic spines using a combined molecular and computational approach.

While actin is extremely well conserved across eukaryotic species and has remained unchanged in the billion years that separate humans and yeast (Gunning et al., 2015), there is another cellular-level memory system even more ancient and conserved. Genes and the mechanisms that regulate gene expression make possible the existence of a diversity of cell subtypes. This so called *epigenetic transcriptional memory system* allows cells to differentiate into specific populations (e.g. neutrocytes, astrocytes, neurons) in response to transient stimuli (Ringrose & Paro, 2004).

Genomes of course determine many physiological features of organisms, including the general makeup of their nervous system, but they don't determine the precious memories for which they become populated. These memories, collected over a lifespan of interacting with the world are precariously stored within a meshwork of living cells, and as such are susceptible to age-related degeneration and degenerative disease. For example Alzheimer's disease (AD) is a chronic neurodegenerative disease that is the most common cause of dementia in humans (Burns & Iliffe, 2009). AD is characterized by a cluster of symptoms, most notably progressive memory loss and cognitive dysfunction (Lambert et al., 2013). This disease currently affects 47 million people worldwide, and with no current treatments is projected to increase threefold by 2050 (Tiwari et al., 2019).

AD is highly heritable (>75% heritability) but genetically complex (Avramopoulos, 2009). The APOE gene (which encodes apolipoprotein E) has been consistently identified as a susceptibility locus for

late-onset Alzheimer's disease (Corder, et al., 1993; Farrer, et al., 1997; Genin, et al., 2011). Since the identification of APOE, genome-wide association studies (GWAS) on AD have recruited substantial participant pools. These sequencing studies have identified an additional 15+ susceptibility loci, including: CLU, CR1, PICALM (Lambert, et al., 2009; Harold, et al., 2009), BIN1 (Seshadri, et al., 2010), MS4A6A, ABCA7, EPHA1, CD33, CD2AP (Hollingworth, et al., 2011; Naj, et al., 2011), TREM2 (Guerreiro, et al., 2013), HLA-DRB5–DRB1, PTK2B, SORL1, SLC24A4-RIN3, and DSG2 (Lambert, et al., 2013). Yet, as you will see in Chapter 1 (Figure-4D) this growing list still provides rather limited information about AD risk.

In general GWAS methods aim to identify single nucleotide polymorphisms (SNPs) with case-control asymmetries surpassing a genome-wide statistical significance ($p < 5 \times 10^{-8}$). In this pursuit there are often SNPs identified to be marginally below this threshold. While some of these marginal SNPs arise due to chance, others fail to reach significance due to low statistical power (particularly when the minor allele is relatively rare). Indeed it has been shown that *polygenic risk models* that incorporate sub-alpha SNPs are significantly better at predicting AD outcomes (Escott-Price, et al., 2015). Chapter 1 of this dissertation presents a novel approach to polygenic risk quantification. Using exome SNPs from approximately 5000 AD cases and 5000 controls, this study is the first to use artificial neural networks to estimate AD risk. Furthermore, we present a method for quantifying a SNP's impact on NN output, which identifies hundreds of novel AD-linked SNPs. Importantly the estimates provided by this method correlated with AD onset age and neuropathology severity (i.e. BRAAK score). Our hope is that these findings will yield tools to assist with AD diagnosis, lead to a better mechanistic understanding of AD, and ultimately preserve precious memories.

WORKS CITED

Albrecht, E., Petty, H. (1998). Cellular memory: Neutrophil orientation reverses during temporally decreasing chemoattractant concentrations. *Proceedings of the National Academy of Sciences* 95(9), 5039-5044. <https://dx.doi.org/10.1073/pnas.95.9.5039>

- Anagnostaras, S., Maren, S., Fanselow, M. (1999). Temporally Graded Retrograde Amnesia of Contextual Fear after Hippocampal Damage in Rats: Within-Subjects Examination. *The Journal of Neuroscience* 19(3), 1106-1114. <https://dx.doi.org/10.1523/jneurosci.19-03-01106.1999>
- Avramopoulos, D. (2009). Genetics of Alzheimer's disease: recent advances. *Genome Medicine* 1(3), 34. <https://dx.doi.org/10.1186/gm34>
- Bosch, M., Castro, J., Saneyoshi, T., Matsuno, H., Sur, M. (2014). Structural and molecular remodeling of dendritic spine substructures during long-term potentiation. *Neuron* 82(2), 444 - 459. <https://dx.doi.org/10.1016/j.neuron.2014.03.021>
- Bourne, J., Harris, K. (2012). Nanoscale analysis of structural synaptic plasticity. *Current Opinion in Neurobiology* 22(3), 372 - 382. <https://dx.doi.org/10.1016/j.conb.2011.10.019>
- Burns, A., Iliffe, S. (2009). Alzheimer's disease. *BMJ (Clinical research ed.)* 338(feb05 1), b158. <https://dx.doi.org/10.1136/bmj.b158>
- Corder, E., Saunders, A., Strittmatter, W., Schmechel, D., Gaskell, P., Small, G., Roses, A., Haines, J., Pericak-Vance, M. (1993). Gene dose of apolipoprotein E type 4 allele and the risk of Alzheimer's disease in late onset families. *Science* 261(5123), 921-923. <https://dx.doi.org/10.1126/science.8346443>
- Escott-Price, V., Myers, A., Huentelman, M., Hardy, J. (2017). Polygenic risk score analysis of pathologically confirmed Alzheimer disease. *Annals of neurology* 82(2), 311 - 314. <https://dx.doi.org/10.1002/ana.24999>
- Farrer, L., Cupples, L., Haines, J., Hyman, B., Kukull, W., Mayeux, R., Myers, R., Pericak-Vance, M., Risch, N., Duijn, C. (1997). Effects of age, sex, and ethnicity on the association between apolipoprotein E genotype and Alzheimer disease. A meta-analysis. APOE and Alzheimer Disease Meta Analysis Consortium. *JAMA* 278(16), 1349 - 1356.
- Fischer, M., Kaech, S., Knutti, D., Matus, A. (1998). Rapid Actin-Based Plasticity in Dendritic Spines. *Neuron* 20(5), 847 - 854. [https://dx.doi.org/10.1016/s0896-6273\(00\)80467-5](https://dx.doi.org/10.1016/s0896-6273(00)80467-5)
- Fukazawa, Y., Saitoh, Y., Ozawa, F., Ohta, Y., Mizuno, K., Inokuchi, K. (2003). Hippocampal LTP Is Accompanied by Enhanced F-Actin Content within the Dendritic Spine that Is Essential for Late LTP Maintenance In Vivo. *Neuron* 38(3), 1 - 14. [https://dx.doi.org/10.1016/s0896-6273\(03\)00206-x](https://dx.doi.org/10.1016/s0896-6273(03)00206-x)
- Genin, E., et al. (2011). APOE and Alzheimer disease: a major gene with semi-dominant inheritance. *Molecular Psychiatry* 16(9), 903-907. <https://dx.doi.org/10.1038/mp.2011.52>

- Gremel, C., Cunningham, C. (2010). Effects of disconnection of amygdala dopamine and nucleus accumbens N-methyl-d-aspartate receptors on ethanol-seeking behavior in mice. *European Journal of Neuroscience* 31(1), 148-155. <https://dx.doi.org/10.1111/j.1460-9568.2009.07044.x>
- Guerreiro, R., Wojtas, A., Bras, J., Carrasquillo, M., Rogaeva, E., Majounie, E., Cruchaga, C., Sassi, C., Kauwe, J., Younkin, S., Hazrati, L., Collinge, J., Pocock, J., Lashley, T., Williams, J., Lambert, J., Amouyel, P., Goate, A., Rademakers, R., Morgan, K., Powell, J., George-Hyslop, P., Singleton, A., Hardy, J. (2013). TREM2 Variants in Alzheimer's Disease. *The New England Journal of Medicine* 368(2), 117-127. <https://dx.doi.org/10.1056/nejmoa1211851>
- Gunning, P., Ghoshdastider, U., Whitaker, S., Popp, D., Robinson, R. (2015). The evolution of compositionally and functionally distinct actin filaments. *Journal of Cell Science* 128(11), 2009 - 2019. <https://dx.doi.org/10.1242/jcs.165563>
- Halpain, S., Hipolito, A., Saffer, L. (1998). Regulation of F-Actin Stability in Dendritic Spines by Glutamate Receptors and Calcineurin. *Journal of Neuroscience* 18(23), 9835-9844. <https://dx.doi.org/10.1523/jneurosci.18-23-09835.1998>
- Harold, D., et al. (2009). Genome-wide association study identifies variants at CLU and PICALM associated with Alzheimer's disease. *Nature Genetics* 41(10), 1088-1093. <https://dx.doi.org/10.1038/ng.440>
- Hollingworth, P., et al. (2011). Common variants at ABCA7, MS4A6A/MS4A4E, EPHA1, CD33 and CD2AP are associated with Alzheimer's disease. *Nature Genetics* 43(5), 429-435. <https://dx.doi.org/10.1038/ng.803>
- Holtmaat, A., Trachtenberg, J., Wilbrecht, L., Shepherd, G., Zhang, X., Knott, G., Svoboda, K. (2005). Transient and Persistent Dendritic Spines in the Neocortex In Vivo. *Neuron* 45(2), 279-291. <https://dx.doi.org/10.1016/j.neuron.2005.01.003>
- Huganir, R., Nicoll, R. (2013). AMPARs and Synaptic Plasticity: The Last 25 Years. *Neuron* 80(3), 704 - 717. <https://dx.doi.org/10.1016/j.neuron.2013.10.025>
- Kessels, H., Malinow, R. (2009). Synaptic AMPA Receptor Plasticity and Behavior. *Neuron* 61(3), 340 - 350. <https://dx.doi.org/10.1016/j.neuron.2009.01.015>
- Kopec, C., Li, B., Wei, W., Boehm, J., Malinow, R. (2006). Glutamate receptor exocytosis and spine enlargement during chemically induced long-term potentiation. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 26(7), 2000 - 2009. <https://dx.doi.org/10.1523/jneurosci.3918-05.2006>
- Lambert, J., et al. (2013). Meta-analysis of 74,046 individuals identifies 11 new susceptibility loci for Alzheimer's disease. *Nature Genetics* 45(12), 1452-1458. <https://dx.doi.org/10.1038/ng.2802>

- Lambert, J., et al. (2009). Genome-wide association study identifies variants at CLU and CR1 associated with Alzheimer's disease. *Nature Genetics* 41(10), 1094-1099. <https://dx.doi.org/10.1038/ng.439>
- Lin, B., Kramar, E., Bi, X., Brucher, F., Gall, C., Lynch, G. (2005). Theta stimulation polymerizes actin in dendritic spines of hippocampus. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 25(8), 2062 - 2069. <https://dx.doi.org/10.1523/jneurosci.4283-04.2005>
- Maiuri, P., Rupprecht, J., Wieser, S., Ruprecht, V., Bénichou, O., Carpi, N., Coppey, M., De Beco, S., Gov, N., Heisenberg, C., Lage Crespo, C., Lautenschlaeger, F., Le Berre, M., Lennon-Dumenil, A., Raab, M., Thiam, H., Piel, M., Sixt, M., Voituriez, R. (2015). Actin Flows Mediate a Universal Coupling between Cell Speed and Cell Persistence. *Cell* 161(2), 374-386. <https://dx.doi.org/10.1016/j.cell.2015.01.056>
- Makino, H., Malinow, R. (2009). AMPA Receptor Incorporation into Synapses during LTP: The Role of Lateral Movement and Exocytosis. *Neuron* 64(3), 381 - 390. <https://dx.doi.org/10.1016/j.neuron.2009.08.035>
- Malenka, R., Nicoll, R. (1999). Long-term potentiation--a decade of progress? *Science* <https://dx.doi.org/10.1126/science.285.5435.1870>
- Malinow, R., Malenka, R. (2002). AMPA Receptor Trafficking and Synaptic Plasticity. *Annual Review of Neuroscience* 25(1), 103-126. <https://dx.doi.org/10.1146/annurev.neuro.25.112701.142758>
- Malinow, R. (2003). AMPA receptor trafficking and long-term potentiation. *Philosophical Transactions of the Royal Society B: Biological Sciences* 358(1432), 707 - 714. <https://dx.doi.org/10.1098/rstb.2002.1233>
- Nabavi, S., Fox, R., Proulx, C., Lin, J., Tsien, R., Malinow, R. (2014). Engineering a memory with LTD and LTP. *Nature* 511(7509), 348 - 352. <https://dx.doi.org/10.1038/nature13294>
- Naj, A., et al. (2011). Common variants at MS4A4/MS4A6E, CD2AP, CD33 and EPHA1 are associated with late-onset Alzheimer's disease *Nature Genetics* 43(5), 436-441. <https://dx.doi.org/10.1038/ng.801>
- Nicoll, R., Roche, K. (2013). Long-term potentiation: peeling the onion. *Neuropharmacology* 74(C), 18 - 22. <https://dx.doi.org/10.1016/j.neuropharm.2013.02.010>
- Okamoto, K., Nagai, T., Miyawaki, A., Hayashi, Y. (2004). Rapid and persistent modulation of actin dynamics regulates postsynaptic reorganization underlying bidirectional plasticity. *Nature Neuroscience* 7(10), 1104 - 1112. <https://dx.doi.org/10.1038/nn1311>
- Pershin, Y., Fontaine, S., Ventra, M. (2009). Memristive model of amoeba learning. *Physical Review E* 80(2), 15 - 6. <https://dx.doi.org/10.1103/physreve.80.021926>

- Prentice-Mott, H., Meroz, Y., Carlson, A., Levine, M., Davidson, M., Irimia, D., Charras, G., Mahadevan, L., Shah, J. (2016). Directional memory arises from long-lived cytoskeletal asymmetries in polarized chemotactic cells. *Proceedings of the National Academy of Sciences* 113(5), 1267-1272. <https://dx.doi.org/10.1073/pnas.1513289113>
- Ringrose, L., Paro, R. (2004). Epigenetic regulation of cellular memory by the polycomb and trithorax group proteins. *Annual Review of Genetics* 38(1), 413-443. <https://dx.doi.org/10.1146/annurev.genet.38.072902.091907>
- Sejnowski, T. (2020). The unreasonable effectiveness of deep learning in artificial intelligence *Proceedings of the National Academy of Sciences* <https://dx.doi.org/10.1073/pnas.1907373117>
- Seshadri, S., et al. (2010). Genome-wide analysis of genetic loci associated with Alzheimer disease. *JAMA* 303(18), 1832-40. <https://dx.doi.org/10.1001/jama.2010.574>
- Skoge, M., Yue, H., Erickstad, M., Bae, A., Levine, H., Groisman, A., Loomis, W., Rappel, W. (2014). Cellular memory in eukaryotic chemotaxis. *Proceedings of the National Academy of Sciences of the United States of America* 111(40), 14448 - 14453. <https://dx.doi.org/10.1073/pnas.1412197111>
- Tiwari, S., Atluri, V., Kaushik, A., Yndart, A., Nair, M. (2019). Alzheimer's disease: pathogenesis, diagnostics, and therapeutics *International Journal of Nanomedicine* 14(), 5541-5554. <https://dx.doi.org/10.2147/ijn.s200490>
- Westendorf, C., Negrete, J., Bae, A., Sandmann, R., Bodenschatz, E., Beta, C. (2013). Actin cytoskeleton of chemotactic amoebae operates close to the onset of oscillations. *Proceedings of the National Academy of Sciences of the United States of America* 110(10), 3853 - 3858. <https://dx.doi.org/10.1073/pnas.1216629110>
- Wixted, J. (2007). Dual-Process Theory and Signal-Detection Theory of Recognition Memory *Psychological Review* 114(1), 152-176. <https://dx.doi.org/10.1037/0033-295x.114.1.152>

CHAPTER 1

A Machine Learning Method to Identify Genetic Variants
Potentially Associated with Alzheimer's Disease

Monk, B., Rajkovic, A., Petrus, S., Rajkovic, A., Gaasterland, T., and Malinow, R.M.

As it appears in

Frontiers in Genetics

2021

Volume 12, Article 647436



A Machine Learning Method to Identify Genetic Variants Potentially Associated With Alzheimer's Disease

Bradley Monk^{1,2}, Andrei Rajkovic³, Semar Petrus⁴, Aleks Rajkovic⁵, Terry Gaasterland⁴ and Roberto Malinow^{1,6*}

¹ Department of Neurosciences, Center for Neural Circuits and Behavior, School of Medicine, University of California, San Diego, San Diego, CA, United States, ² Cognitive Science & Psychology IDP, University of California, San Diego, San Diego, CA, United States, ³ Department of Computer Science, Royal Holloway, University of London, Egham, United Kingdom, ⁴ Institute for Genomic Medicine, Scripps Institution of Oceanography, University of California, San Diego, San Diego, CA, United States, ⁵ Department of Pathology, Department of Obstetrics, Gynecology and Reproductive Sciences, University of California, San Francisco, San Francisco, CA, United States, ⁶ Section of Neurobiology, Division of Biological Sciences, University of California, San Diego, San Diego, CA, United States

OPEN ACCESS

Edited by:

Yun Xiao,
Harbin Medical University, China

Reviewed by:

Xiaodong Zheng,
Anhui Medical University, China
Michelle Lupton,
QIMR Berghofer Medical Research
Institute, Australia

*Correspondence:

Roberto Malinow
rmalinow@ucsd.edu

Specialty section:

This article was submitted to
Computational Genomics,
a section of the journal
Frontiers in Genetics

Received: 29 December 2020

Accepted: 06 April 2021

Published: 14 June 2021

Citation:

Monk B, Rajkovic A, Petrus S,
Rajkovic A, Gaasterland T and
Malinow R (2021) A Machine Learning
Method to Identify Genetic Variants
Potentially Associated With
Alzheimer's Disease.
Front. Genet. 12:647436.
doi: 10.3389/fgene.2021.647436

There is hope that genomic information will assist prediction, treatment, and understanding of Alzheimer's disease (AD). Here, using exome data from ~10,000 individuals, we explore machine learning neural network (NN) methods to estimate the impact of SNPs (i.e., genetic variants) on AD risk. We develop an NN-based method (netSNP) that identifies hundreds of novel potentially protective or at-risk AD-associated SNPs (along with an effect measure); the majority with frequency under 0.01. For case individuals, the number of "protective" (or "at-risk") netSNP-identified SNPs in their genome correlates positively (or inversely) with their age of AD diagnosis and inversely (or positively) with autopsy neuropathology. The effect measure increases correlations. Simulations suggest our results are not due to genetic linkage, overfitting, or bias introduced by netSNP. These findings suggest that netSNP can identify SNPs associated with AD pathophysiology that may assist with the diagnosis and mechanistic understanding of the disease.

Keywords: machine learning, neural network, Alzheimer's, disease, polygenic

INTRODUCTION

Alzheimer's disease (AD), the most common form of dementia, is heritable [58–79%, estimated from twin studies (Gatz et al., 2006)], and highly polygenic (Cauwenberghe et al., 2015). Mutations in three genes (*APP*, *PS1*, *PS2*) cause rare forms of the disease [accounting for ~1% of AD cases (Mendez, 2017)], which shows autosomal dominant transmission with high penetrance and displays an early onset [generally before age 60 (Carmona et al., 2018)]. In the more common form of the disease, late onset AD (LOAD), *APOE* has been established unequivocally as a susceptibility gene (Saunders et al., 1993) with several dozen other genetic loci receiving genetic support (Carmona et al., 2018; Jansen et al., 2019; Kunkle et al., 2019).

The neuropathology of AD is defined by the presence of extracellular senile plaques containing amyloid beta 42 and intracellular neurofibrillary tangles containing hyperphosphorylated tau protein (DeTure and Dickson, 2019). The neuropathological progression of disease has been best described using the Braak staging scheme (I–VI) (Braak et al., 2006). The most important genetic

variant in LOAD is the *APOE* $\epsilon 4$ isoform, which predisposes patients to an earlier appearance of AD and a higher Braak score. The role of *APOE* or other identified genetic variants in the pathophysiology of AD is not well understood (Sisodia and George-Hyslop, 2002; Koffie et al., 2012; Karch and Goate, 2014; Shi et al., 2017). Currently available disease biomarkers can be expensive, labor intensive, and do not provide a definitive clinical diagnosis (Gustaw-Rothenberg et al., 2010; Hampel et al., 2018; Jack et al., 2018; Penner et al., 2019). The identification of additional LOAD-linked genetic variants could potentially increase diagnostic accuracy, increase our understanding of the disease, and unmask potential drug targets.

In 2009 two high-powered genome-wide association (GWAS) studies were published that identified, along with *APOE*, several single nucleotide polymorphism (SNPs) loci significantly linked to AD, including SNPs in PICALM, CLU, and CR1 (Harold et al., 2009; Lambert et al., 2009). To continue the search for genetic targets linked to AD, the Alzheimer's Disease Sequencing Project (ADSP) was established as a collaboration between the National Human Genome Research Institute (NHGRI) and the National Institute on Aging (NIA). As part of this effort, whole-exome sequencing was performed on 5,740 LOAD cases and 5,096 cognitively normal, older individuals, serving as controls (Bis et al., 2018). The overarching goals of this initiative have been to identify novel genomic targets that contribute risk or confer protection toward AD outcomes, and to develop new insights as to why some at-risk individuals do not develop AD. Indeed, data from this project have been used to identify a number of novel SNPs linked to AD (Beecham et al., 2018; Bis et al., 2018; Raghavan et al., 2018; Ma et al., 2019; Patel et al., 2019; Zhang et al., 2019).

Recent studies have presented polygenic risk score (PRS) models for estimating AD risk (Escott-Price et al., 2015, 2017; Desikan et al., 2017; Zhang et al., 2020). In these models, GWAS summary data were used to identify AD-linked genomic variants and to assign each variant a coefficient based on their case-control asymmetries. While PRS models are a powerful method to identify individuals at risk for a disease, we believe they could provide another powerful utility – identifying novel genetic variants that confer AD risk or protection that escape GWAS identification for a number of factors, including rarity and potential interactions (linear and non-linear) with other variants. However, in order to capture these interactions the PRS model needs to be based on individual genotypes (not just GWAS summary data). Here we developed such a model based on individual AD case and control SNP data provided by ADSP. Specifically, artificial neural networks (Sejnowski, 2020) were trained using individual case and control genotypes to estimate polygenic risk. A primary aim of this study was to develop a machine learning-based method (netSNP) that can be used to identify the importance of individual SNPs in a complex polygenic classifier's decision making process.

netSNP can identify hundreds of AD-linked SNPs, most of which have a low population frequency (<0.01). Supporting the validity of our method are the observations that the number of AD-linked SNPs identified by this method that are harbored by an individual diagnosed with AD is correlated with the age

at which that individual's AD was diagnosed as well as their brain pathology. In particular, the number of risk- (or protection-) linked SNPs an individual harbors correlates negatively (or positively) with the age at which an individual is diagnosed with AD and with their Braak score (i.e., individuals with more risk SNPs had AD at earlier ages and higher Braak scores; individuals with more protective SNPs had AD at later ages and lower Braak scores). Furthermore, scaling the SNPs with a netSNP-derived "importance factor" further increases the correlations. Thus, these correlations provide support for the view that this method correctly identifies AD-linked SNPs and correctly quantifies their importance.

RESULTS

Dataset Pipeline, Case:Control Balancing, and SNP Properties

A large variant call format (VCF) datafile [~ 200 GB; Alzheimer's Disease Sequencing Project, ADSP (Beecham et al., 2017)] containing SNP information (i.e., reference or alternate allele for ~ 1.4 million SNP sites) on 11,000 individuals over the age of 60 (Northern European descent; 6,000 diagnosed with AD, and 5,000 aged non-AD controls), was organized into a more manageable file (~ 2 GB; N.B.: a VCF datafile contains mainly zeros – indicating reference alleles – since $>95\%$ of minor allele frequencies are <0.01) permitting rapid queries regarding SNP content for any individual (see section "Materials and Methods"). The minor frequency allele (MFA) and reference allele count were determined at each SNP locus, separately for case and control groups. The Fisher's exact test was used to quantify the probability (FishP) that the observed case/control minor allele asymmetry could be due to chance.

The ADSP dataset consists of SNP information originating from 24 cohort groups (Beecham et al., 2017; Crane et al., 2017; Naj et al., 2018). We initially tested if a neural net (NN) classifier could be trained (Moller, 1993), with SNPs as features (50 SNPs with the lowest FishP values; 50 features were chosen as this was computationally tractable; see section "Materials and Methods"), to identify from which cohort an individual originated. Indeed, the classifier could identify cohort identity for each individual with 50% accuracy, much above the 4% expected by chance (Supplementary Figures 1, 2). This was of concern, because given the large case:control imbalance in many cohorts (see Supplementary Figure 2), the classifier could use cohort information to help identify patient AD status. Thus, the SNPs would indicate something about the cohort (e.g., platform- or probe-specific aspects of cohorts) rather than the disease. To avoid this potentially confounding issue, we balanced cohorts. In short, (a) only cohorts with at least 20% of the cohort being cases or controls were used; and (b) the same number of case and control individuals from each cohort was used in training sets (see Supplementary Figure 2 and "Materials and Methods" for details).

A quantile-quantile (Q-Q) plot of the $-\log(\text{FishP})$ values of a balanced dataset plotted against a similar computation of the same dataset with shuffled case-control labels shows that

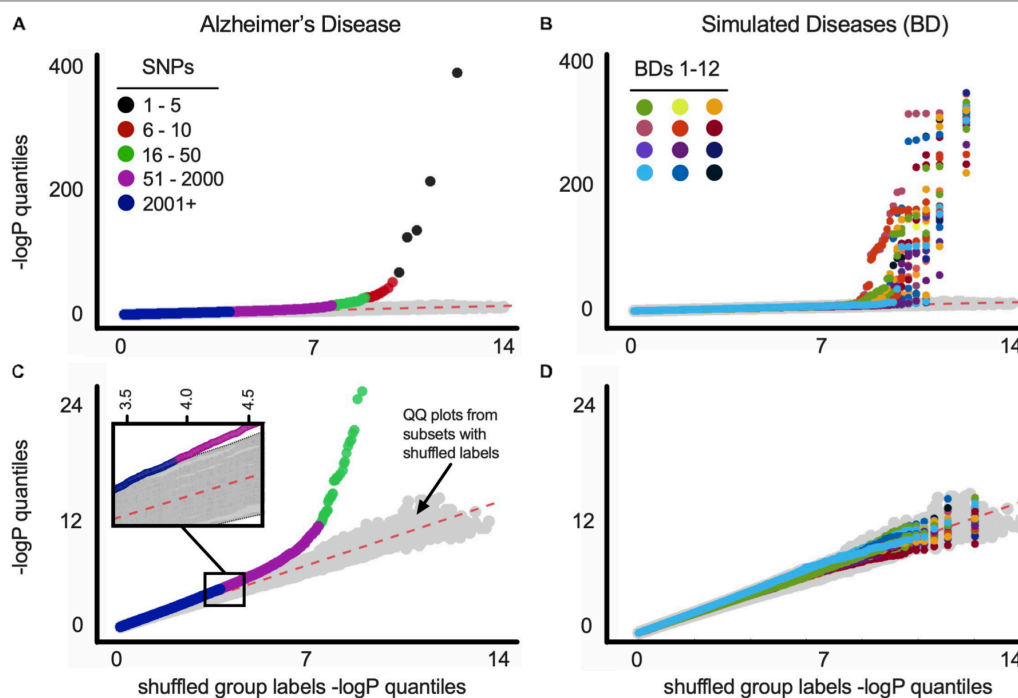


FIGURE 1 | (A) Q-Q plots of balanced ADSP dataset for Alzheimer's disease. Gray symbols (here and below) represent Q-Q plot of 100 random versus random distributions (i.e., chance). See text for details. SNP order [based on $-\log(P)$] is indicated by colors (see legend). (B) Q-Q plots of balanced ADSP dataset for 12 constructed (simulated) diseases (BDs). Each BD is represented by a different color. (C) Same as above with SNPs from *APOE*-residing chromosome 19 removed before p -value quantiles were computed. Inset is a blow up of the indicated region, showing magenta SNPs fall outside the 100 random versus random distributions (gray symbols). SNP order [based on $-\log(P)$] is indicated by colors (see legend). (D) Same as above with SNPs from *BDgene*-residing chromosome removed before p -value quantiles were computed. Plot as in B is shown for all BD1-33 in Supplementary Figure 12.

most of the case-control minor allele asymmetries across the 1.4 million SNP loci can be explained by chance (i.e., lie on the $x = y$ line; **Figures 1A–D**). For comparison we plotted 100 Q-Q plots, wherein $-\log(\text{FishP})$ values from one shuffled dataset was plotted against $-\log(\text{FishP})$ values from another shuffled dataset (**Figures 1A–D**, gray symbols). For the AD population, in a few SNPs from *APOE* and (its physically close linkage partner) *TOMM40* genes (Yu et al., 2007; Guerreiro and Hardy, 2012), the observed distribution of reference allele (Ref) and MFA in the case and control populations was far (orders of magnitude) from what can be accounted for by chance (**Figures 1A,C**).

To address the possibility that artifacts can account for SNPs off the $x = y$ line (e.g., SNPs being linked to *APOE*, rather than to AD), we constructed 33 separate simulated diseases (“bad diseases,” BDs) using all ADSP individuals (cf., Bulik-Sullivan et al., 2015). Each BD was based on an existing gene (*BDgene*) that has two SNPs with frequencies in our population very close to *APOE* ϵ 4 (0.147, E4-like) and *APOE* ϵ 2 (0.076, E2-like); see **Table 1**, MAF (minor allele frequency) columns. Individuals with the *BDgene* genotype (i.e., having E4-like or E2-like SNPs) in the ADSP dataset were randomly ascribed to have BDs based on control/case odds ratio of *APOE* ϵ 4 (0.30) and *APOE* ϵ 2 (2.41) for AD. Individuals without *BDgene* SNPs were randomly assigned

based on the odds ratio of individuals without *APOE* ϵ 4 or ϵ 2 (i.e., *APOE* ϵ 33 = 0.89). FishP values were computed for each SNP from true (AD) and simulated (BDs) diseases from balanced data sets, and Q-Q plots were generated (**Figures 1B,D**). Plots including all SNPs showed many with FishP values outside what could be accounted for by chance for both AD and BDs (**Figures 1A,B**). However, if SNPs from chromosome 19 (where *APOE* resides) or the chromosome with *BDgene* were removed, only SNPs for AD could not be accounted for by chance (**Figures 1C,D**). This result is consistent with the view previously observed that AD is a highly polygenic disorder (Cauwenbergh et al., 2015; Escott-Price et al., 2017) as there was a considerable asymmetry in MAF between case and control populations for over 2,000 SNPs (see **Figure 1D**). While artifacts related to data stratification can account for this behavior in Q-Q plots (Lander and Schork, 1994; Slatkin, 2007), cohort balancing and our simulations argue against such artifacts for our dataset, and support the existence of a large number of SNPs linked to AD, consistent with previous results (Cauwenbergh et al., 2015; Escott-Price et al., 2017).

NN Construction and Performance

Once the cohorts were balanced, we calculated the FishP values for SNPs from a “training set” composed of 3,200 randomly

TABLE 1 | netSNP identified tSNPs with greatest absolute average CVt when APOE locus variants were not excluded from the training set.

tSNPs predicted to confer most protection against AD						tSNPs predicted to confer the most risk for AD					
Chr	Pos	Gene	<i>mCVt</i>	FishP	MAF	Chr	Pos	Gene	<i>mCVt</i>	FishP	MAF
5	612,536	CEP72	-0.243	5.7E-03	0.002	11	10,327,875	ADM	0.289	4.1E-08	0.008
6	1,390,303	FOXF2	-0.182	1.4E-02	0.003	19	45,411,941	APOE $\epsilon 4^*$	0.261	3.4E-111	0.135
4	110,638,764	PLA2G12A	-0.178	2.8E-02	0.002	7	23,213,734	KLHL7	0.217	5.1E-03	0.003
1	16,890,642	NBPF1	-0.174	3.3E-02	0.002	9	130,439,029	STXBP1	0.207	1.9E-02	0.001
11	1,017,294	MUC6	-0.157	1.4E-07	0.01	20	37,258,198	ARHGAP40	0.203	1.1E-02	0.002
1	40,961,395	ZFP69	-0.156	4.8E-04	0.002	15	41,862,356	TYRO3	0.197	8.4E-18	0.018
15	50,154,563	ATP8B4	-0.156	2.3E-02	0.004	6	146,276,263	SHPRH	0.195	6.6E-03	0.002
19	52,793,834	ZNF766	-0.155	2.8E-02	0.002	1	228,879,367	RHOU	0.195	6.4E-03	0.004
15	64,017,685	HERC1	-0.152	3.5E-03	0.004	9	131,398,647	WDR34	0.19	7.0E-03	0.002
19	45,412,079	APOE $\epsilon 2^*$	-0.152	7.1E-38	0.079	19	52,497,235	ZNF615	0.188	3.3E-02	0.003
16	8,740,006	METTL22	-0.15	2.1E-03	0.002	12	85,450,243	LRRIQ1	0.188	1.0E-02	0.006
9	139,396,933	NOTCH1	-0.143	3.3E-02	0.004	15	25,963,545	ATP10A	0.185	1.2E-02	0.002
19	18,561,473	ELL	-0.137	7.3E-03	0.008	12	108,011,971	BTBD11	0.183	3.8E-03	0.007
11	57,467,411	ZDHC5	-0.133	3.5E-03	0.002	9	107,533,232	NIPSNAP3B	0.181	1.1E-02	0.003
9	100,372,648	TSTD2	-0.131	2.2E-02	0.003	1	8,420,270	RERE	0.18	3.0E-02	0.004
1	65,120,426	CACHD1	-0.131	1.0E-02	0.002	8	10,480,495	RP1L1	0.179	3.1E-02	0.003
12	69,113,184	NUP107	-0.126	5.6E-03	0.006	4	5,682,993	EVC2	0.178	2.2E-02	0.004
5	145,508,644	LARS	-0.126	1.2E-02	0.006	5	140,530,973	PCDHB6	0.178	1.4E-02	0.002
7	6,561,105	GRID2IP	-0.125	2.9E-04	0.002	6	30,712,298	IER3	0.177	3.3E-03	0.007
19	43,268,061	PSG8	-0.125	2.9E-02	0.004	15	50,264,839	ATP8B4	0.176	1.1E-02	0.008
11	47,264,353	ACP2	-0.125	1.4E-03	0.004	16	3,604,305	NLRC3	0.176	2.7E-02	0.002
6	7,405,242	RIOK1	-0.124	1.5E-02	0.003	22	46,780,446	CELSR1	0.174	1.9E-02	0.003
3	146,167,089	PLSCR2	-0.123	2.3E-02	0.003	19	39,103,307	MAP4K1	0.173	1.9E-02	0.001
16	30,775,522	RNF40	-0.123	2.9E-02	0.006	1	89,579,827	GBP2	0.173	2.5E-02	0.005
9	139,008,644	C9orf69	-0.121	2.9E-03	0.003	12	50,500,080	GPD1	0.173	2.6E-03	0.002

Rows 26:1000 available online

Rows 26:1000 available online

* Previously published AD-linked gene

* Previously published AD-linked gene

chosen individuals (case & controls; equal number of each) and used the 50 SNPs with the lowest FishP values in the training set to train an NN classifier to predict if an individual was a case or control (Figure 2, left; see section “Materials and Methods”). Briefly (Demuth et al., 2014), an artificial NN was trained to classify cases vs. controls using genotypes (for 50 SNPs) of individuals in the training set. The NN was initialized with random weights connecting each node, so the initial prediction y was random (each y was a real number scaled between 0.5, the control label, and 0.5, the case label). This prediction, also known as a classifier value (CV), was evaluated against the true label (case or control) using a loss function, and the network weights were updated using an optimization function. Throughout training the optimizer adjusts NN weights, working to minimize the loss function. Training concluded when the NN weights were considered optimal (within the constraints of the stopping criteria and cross validation; see section “Materials and Methods”), at which point the NN weights remain fixed. Thus, additional input to the NN would yield CV predictions, but would not change network weights or alter the model in any regard.

After the NN was trained as described above, it was applied to the 50 SNPs of each individual from the “holdout set”

(1,500 individuals randomly chosen who were not included in the training set), providing a CV for each. Overall these CVs correlated well with actual AD status of each individual (case, red; control, blue; Figure 2A and Supplementary Figure 4). Using a classification threshold of zero (such that any positive CV was predicted as case, and any negative CV was predicted as control), the classifier accuracy was 67.3% (SD = 0.3%, see section “Materials and Methods”). The NN performance with 50 SNPs was significantly better than what could be achieved using only SNPs from the APOE locus (62.2%). It also performed better than a logistic regression model using the same 50 SNPs (64.2%, $p < 10e-20$, McNemar Test). When only considering individuals with CVs closer to 0.5 or 0.5, the accuracy of the NN increased. For individuals with CVs in the outer quartiles, prediction accuracy was 76.4% (SD = 0.5%); for those with a CV ranked in the upper 12.5% and lower 12.5% quantiles, the classification accuracy was 82.6% (SD = 0.6%) (see Figure 2C).

We next trained an NN using a set of 50 SNPs (a) not containing APOE gene SNPs, or (b) not containing the 22 previously published AD-associated SNPs (Carmona et al., 2018), or (c) with the 51–100 lowest FishP values. The resulting accuracy and receiver operator characteristic (ROC) curves, which provide a measure of the sensitivity and specificity of a

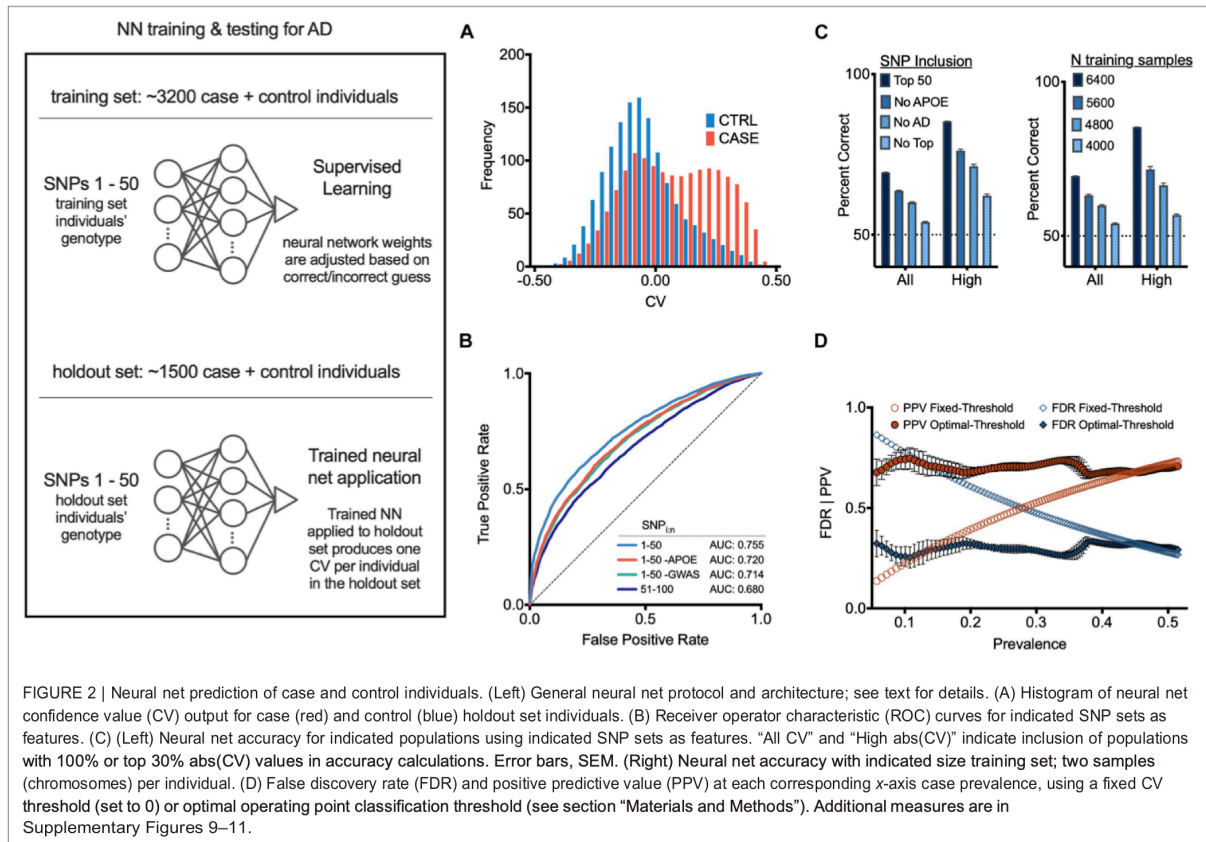


FIGURE 2 | Neural net prediction of case and control individuals. (Left) General neural net protocol and architecture; see text for details. (A) Histogram of neural net confidence value (CV) output for case (red) and control (blue) holdout set individuals. (B) Receiver operator characteristic (ROC) curves for indicated SNP sets as features. (C) (Left) Neural net accuracy for indicated populations using indicated SNP sets as features. "All CV" and "High abs(CV)" indicate inclusion of populations with 100% or top 30% abs(CV) values in accuracy calculations. Error bars, SEM. (Right) Neural net accuracy with indicated size training set; two samples (chromosomes) per individual. (D) False discovery rate (FDR) and positive predictive value (PPV) at each corresponding x-axis case prevalence, using a fixed CV threshold (set to 0) or optimal operating point classification threshold (see section "Materials and Methods"). Additional measures are in Supplementary Figures 9–11.

method (Koen et al., 2016), were all above chance in predicting AD status of an individual (Figures 2B,C). Reducing the size of the training set reduced the accuracy in a roughly linear fashion (Figure 2C), suggesting that the NN accuracy did not asymptote at 3,200 individuals, and that gathering SNP information from more individuals would increase NN accuracy. The area under the ROC curve (AUC) for our NN model with 50 SNPs was 0.755. Further analysis of NN hyperparameters such as the number of SNPs, which SNPs were employed, NN architecture, etc., may improve NN performance; we note that producing an optimal NN was not the primary goal of this study. Other methods, such as PCA (Jolliffe, 1986; Selzam et al., 2018), or Random Forest (Goldstein et al., 2011) analyses were not examined.

Due to the cohort counterbalancing requirement (see above), the prevalence of AD in our training set was 0.5. Since disease prevalence in most populations will almost certainly be lower than 0.5, we quantified signal detection metrics for a range of disease prevalence rates from 0.05 to 0.5 (0.05 is the approximate AD prevalence at age 75), using the optimal operating point (OOP) for each respective base rate (see section "Materials and Methods"). Using the OOP, the false discovery rate was largely independent of prevalence for values from 0.05 to 0.5 (Figure 2D). Similarly, the same optimal threshold maintained a largely constant positive predictive rate (Figure 2D). Thus,

computing an NN with training data composed of an equal number of cases and controls can be used despite a low disease prevalence.

netSNP Description and Application

While NNs can perform well in solving complex problems, determining the importance of different NN input features (in this case, different SNPs) is difficult to assess. With this in mind, we developed a method (netSNP) using a modification of the standard NN protocol, aimed to assess the impact of any SNP on conferring AD risk or protection. Specifically, we derived a quantitative measure for the impact of an SNP on the output of an NN.

netSNP is a modification of the *Permutation Importance* method used in machine learning (Altmann et al., 2010; Molnar, 2021), which we have adapted for use with polygenic models. In general Permutation Importance is used to address the question "What variables have the biggest impact on the predictions of a trained neural network classifier?" Permutation Importance computations are performed after a model has already been fitted, and works using a basic strategy: a single predictor variable is modified in the input data, leaving all the other predictor variables unchanged, and examining how this affects classifier

performance. This procedure is then repeated, one variable at a time, for all the predictor variables used in the model. This permits one to determine the relative effect of each predictor variable. The netSNP method uses a similar strategy. For a specific SNP, netSNP addresses this question “if this SNP is artificially made homozygous for the MFA, what impact does it have on the classifier output?” If the average CV shifts to the right (e.g., goes from 0.1 to 0.3) when a SNP is set to homozygous for the MFA, netSNP deems this SNP to confer risk. If the average CV shifts to the left (e.g., goes from 0.1 to 0.2) when a SNP is set to homozygous for the MFA, netSNP deems this SNP to confer protection.

To demonstrate the netSNP method on a specific example, we used netSNP to compute the impact of the *APOE* genotype on AD risk. From a balanced dataset, we randomly chose a training set composed of 3,200 cases and controls (which contained individuals with all *APOE* genomic variants; i.e., *APOE* ϵ 22, ϵ 23, ϵ 24, ϵ 33, ϵ 34, and ϵ 44). This set was used to train an NN (which we call NN ϵ) to identify cases or controls based on their top 50 SNPs (see section “Materials and Methods” and **Figure 3** left panel, top). After this training session, NN ϵ was not modified in the subsequent analysis of the *APOE* genomic variants. We then applied NN ϵ to a holdout set of 1,500 individuals (**Figure 3**, left panel, bottom), producing 1,500 CV outputs with a distribution shown in **Figure 3A** (dashed line; this is used as a *baseline* for comparisons). We next reasoned that the impact of a specific *APOE* genotype on NN ϵ predictions could be assessed by artificially modifying the *APOE* genotype of every holdout set individual to that specific *APOE* genotype. For instance, to assess the impact of the ϵ 22 genotype, we artificially assigned every holdout set individual the *APOE* ϵ 22 genotype (keeping non-*APOE* genotypes of each individual unaltered). After applying NN ϵ to these modified genotypes, the distribution of CV outputs was strongly shifted leftward compared to the *baseline* distribution (**Figure 3A**, compare blue distribution to dashed line). Alternatively, if we assigned all holdout set individuals the ϵ 44 genotype, the CV distribution shifted significantly rightward from *baseline* (**Figure 3A**, compare orange distribution to dashed line). Falling between the ϵ 22 and ϵ 44 distributions were the CV distributions when NN ϵ was applied to holdout set individuals assigned either the ϵ 23, or ϵ 33, or ϵ 24, or ϵ 34 genotype (**Figure 3A**).

We next performed a critical test of netSNP: to determine if the above (i.e., the colored CV distributions in **Figure 3A**) corresponded to distributions when NN ϵ was applied for individuals who did have distinct *APOE* genotypes. To test for this, we created holdout sets with individuals with only one *APOE* genotype (i.e., one holdout set included only *APOE* ϵ 22 individuals, another holdout set only *APOE* ϵ 23 individuals, etc.). We then used NN ϵ to compute CVs for individuals in each of these holdout sets (using the true genotypes for each individual, for 50 SNPs). The resulting CV distribution for true *APOE* genotype holdout sets moved from left to right as *APOE* changed from ϵ 22 to ϵ 44 (**Figure 3B**), closely matching the CV distributions from above, where *APOE* status was assigned to all individuals in the holdout set (compare **Figures 3A,B**). This result suggests that netSNP

can accurately assess the impact of individual SNPs on a classifier output.

Since *APOE* SNPs are known to significantly impact AD risk, this result also suggests that the netSNP method could be used to estimate the impact of many different target SNPs of interest (which we call *tSNPs*) on AD risk. We achieved this for each *tSNP* by performing the following procedure (analogous to the procedure used to test the impact of *APOE* genotypes above; see **Figure 3**, left panel): From a balanced dataset, we randomly chose a training set composed of 3,200 cases and controls. This set was used to train an NN (which we call NN t) to identify cases or controls based on their true genotypes for 50 SNPs (the top 49 SNPs based on FishP value, and the *tSNP* of interest). Then we constructed a holdout set of 1,500 individuals, and applied NN t on each individual, using the same 50 SNPs used in training, and using the true genotypes of each individual. This produced 1,500 baseline CV values. Finally we constructed a holdout set of 1,500 individuals, and used the same 50 SNPs, using their true genotypes for each individual for 49 SNPs, but the *tSNP* was set to be homozygous for the *tSNP* MFA. We then applied the NN t producing 1,500 CV values (which we call *CV t*) which can be plotted in a frequency distribution (**Figure 3**, bottom). We repeated this procedure for many *tSNPs* (see section “Materials and Methods”; **Figure 3C** shows *CV t* distributions for several *tSNPs*). Intuitively, we reasoned that if an SNP had an effect on AD risk, then when evaluated as a *tSNP*, the *CV t* distribution would be shifted compared to the baseline CV distribution – shifts to the left would indicate the MFA SNP is AD-protective (**Figure 3**, left panel, “Protective SNP *t*” distribution); shifts to the right would indicate the MFA SNP incurs AD risk; the larger the shift, the greater the impact on AD. We test this proposal below.

We used netSNP to test 4,000 individual SNPs as *tSNPs*; we chose those SNPs with the 4,000 lowest FishP values. Each *tSNP* was evaluated 20 times (see section “Materials and Methods”) from which a mean *CV t* (*mCV t*) is computed over all holdout set individuals for all 20 runs. Evaluating *APOE* ϵ 4 as a *tSNP* with netSNP resulted in a *CV ϵ 4* distribution that was shifted to the right (**Figure 3C**, 2nd from top; same as **Figure 3A**, green), as expected. Surprisingly, the MFA of an adrenomedullin (*ADM*) SNP shifted the CV distribution more to the right than *APOE* ϵ 4 (*mCV ϵ 4* = 0.26 \pm 0.001; *mCV $_{ADM}$* = 0.29 \pm 0.001). \pm Also, a number of SNPs shifted NN t output CVs more to the left than *APOE* ϵ 2 (e.g., *mCV ϵ 2* = -0.15 \pm 0.001; *mCV $_{CEP72}$* = 0.24 \pm 0.001; see **Table 1** for *tSNPs* with the most extreme *mCV t*). Thus netSNP appears to identify a number of SNPs that can considerably shift NN output CV, potentially identifying SNPs that confer AD protection (shifting CV to the left) and AD risk (shifting CV to the right).

To exclude the artifactual possibility that netSNP was dependent on *APOE*, we repeated the netSNP method with *APOE* (and *TOMM40*) SNPs excluded from the 49 SNPs with the lowest FishP values as features in training NN t (although *APOE* was tested as a target *tSNP*). Results were very similar to the above, with hundreds of *tSNPs* shifting CV to the right (potentially AD risk SNPs) and hundreds of *tSNPs* shifting CV to the left (potentially AD protective SNPs; **Supplementary Table 1**).

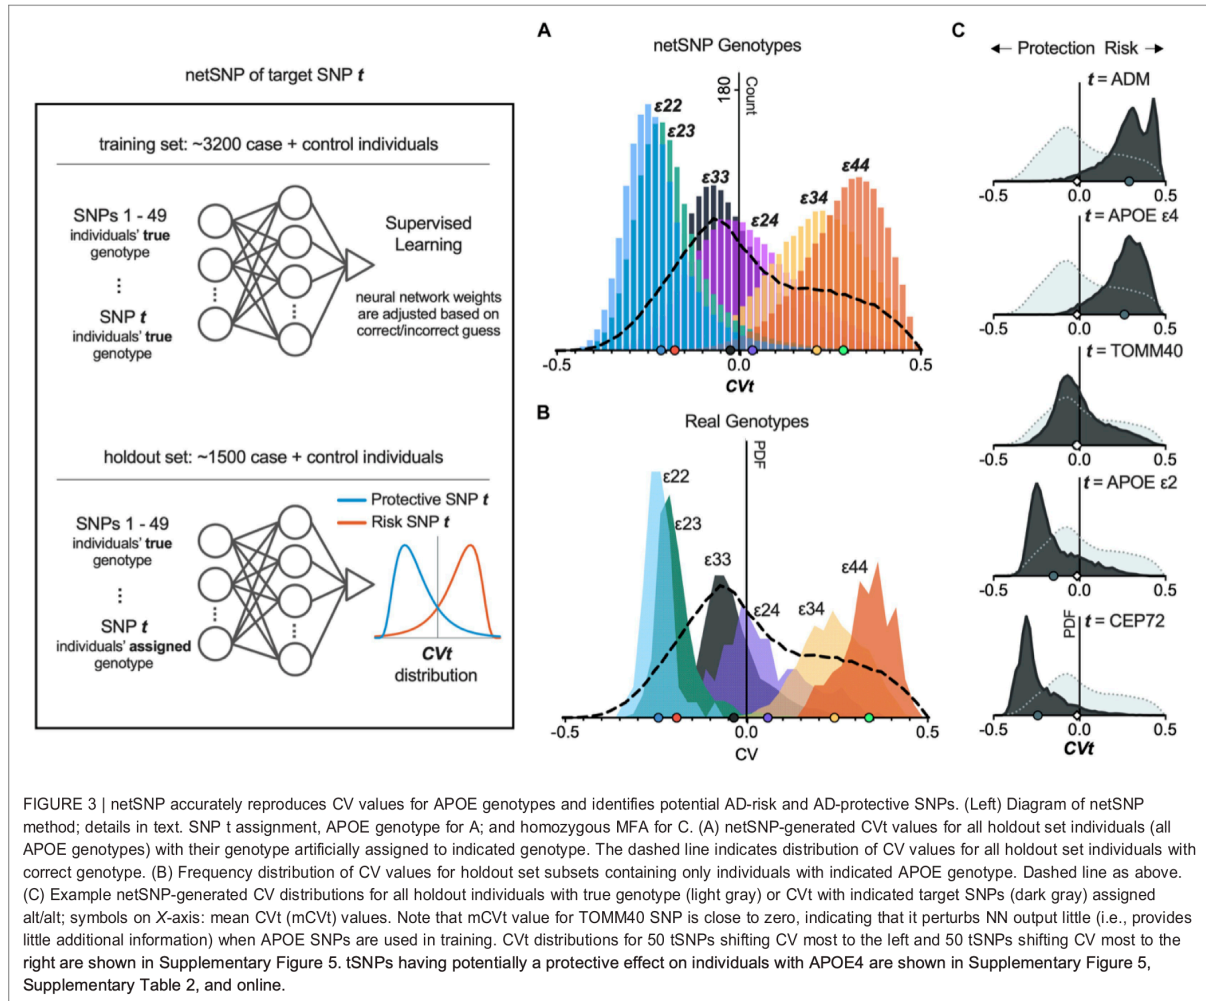


FIGURE 3 | netSNP accurately reproduces CV values for APOE genotypes and identifies potential AD-risk and AD-protective SNPs. (Left) Diagram of netSNP method; details in text. SNP t assignment, APOE genotype for A; and homozygous MFA for C. (A) netSNP-generated CVt values for all holdout set individuals (all APOE genotypes) with their genotype artificially assigned to indicated genotype. The dashed line indicates distribution of CV values for all holdout set individuals with correct genotype. (B) Frequency distribution of CV values for holdout set subsets containing only individuals with indicated APOE genotype. Dashed line as above. (C) Example netSNP-generated CV distributions for all holdout individuals with true genotype (light gray) or CVt with indicated target SNPs assigned alt/alt; symbols on X-axis: mean CVt (mCVt) values. Note that mCVt value for TOMM40 SNP is close to zero, indicating that it perturbs NN output little (i.e., provides little additional information) when APOE SNPs are used in training. CVt distributions for 50 tSNPs shifting CV most to the left and 50 tSNPs shifting CV most to the right are shown in Supplementary Figure 5. tSNPs having potentially a protective effect on individuals with APOE4 are shown in Supplementary Figure 5, Supplementary Table 2, and online.

In general, this method provides a quantitative measure of the impact (as indicated by **mCVt** values) of specific SNPs on NN output, and potentially (see below) the effect of such SNPs on developing AD.

NN and CV as Predictors of AD and Its Pathophysiology

While CV values (computed with or without APOE as an NN feature) predict well the likelihood of an individual being diagnosed with AD (Supplementary Figure 4), we aimed to determine if CV values correlate with the pathophysiology underlying AD. We reasoned that individuals diagnosed with AD at an earlier age may have a more aggressive form of the disease, which could be a consequence of their genetics, and this might be detected by more positive CV values; equivalently, AD diagnosis at an older age may correlate with less aggressive AD pathophysiology, and may have more negative CV values. This reasoning is supported by previous findings with APOE

genotypes (Corder et al., 1993), which we found to also be true in our dataset (Figure 4A). Linear regression fitting shows that, for case individuals, as their APOE ϵ_2 allele count increases, so does their observed disease onset age [$F(2,4750) = 86, \beta = 3.8, p < 2.6e-20$; general linear model, see section “Materials and Methods”]; conversely, the number of APOE ϵ_4 alleles reduces the age of AD diagnosis [$F(2,4750) = 1910, \beta = 8.4, p < 1e-300$]. With this reasoning in mind, we tested and found that the age at which cases were diagnosed with AD could be predicted by their CV values [as computed in section “NN Construction and Performance”]; more positive CV for younger age of AD diagnosis, $F(2,4752) = 571, \beta = 27, p < 2.3e-119$; Figure 4B]. Furthermore, an individual’s CV was positively correlated with Braak score, for case individuals receiving autopsy [$F(2,2025) = 154, \beta = 1.8, p < 4.1e-34$; Figure 4G]. These effects were also highly significant if APOE was not included in the NN calculation of CV [CV vs. age, $F(2,4752) = 422, \beta = -22.4,$

$p = 5.3e-90$; CV vs. Braak, $F(2,2025) = 59$, $\beta = 1.7$, $p = 1.8e-14$]. These findings support the view that the NN output value CV, as described above in section “NN Construction and Performance,” is related to the pathophysiology of AD.

netSNP as Predictor of AD-Linked *tSNPs* and AD Pathophysiology

We next tested if netSNP can identify AD-linked SNPs and can quantify their impact on the likelihood of developing AD. We considered a set of *tSNPs* for which their computed *mCVt* values were significantly ($p < 0.05$) outside the range of *mCVt* values generated by randomly choosing target SNPs from the set of all $1.4 \cdot 10^6$ ADSP SNPs (see section “Materials and Methods”). This resulted in 851 *tSNPs* with *mCVt* < 0 (provisionally indicated “AD-protective *tSNPs*”) and 672 *tSNPs* with *mCVt* > 0 (“AD-risk *tSNPs*”), the majority (64%) with MAF under 0.01. Only some of the previously published AD-linked SNPs (which we exclude from the subsequent validation analysis) are in these sets (see **Table 1**). Using a general linear model, we found that the number of “AD-protective *tSNPs*” harbored by each case individual correlated positively with their age of AD diagnosis [$F(2,4750) = 13.9$, $\beta = 0.072$, $p < 1.9e-4$; **Figure 4C**], while the number of “AD-risk *tSNPs*” they harbored correlated inversely with age of AD diagnosis [$F(2,4750) = 18.2$, $\beta = 0.11$, $p < 1.9e-05$; **Figure 4C**]. Providing *tSNPs* with a *CVt* weight increased the positive correlation between *CVt*-weighted “AD protective *tSNPs*” [$F(2,4750) = 400$, $\beta = 22$, $p < 1.6e-85$], or the negative correlation between *CVt*-weighted “AD risk *tSNPs*” [$F(2,4750) = 404$, $\beta = 25$, $p < 2.4e-86$; - **Figure 4D**] and age of AD diagnosis. Interestingly, the number of previously published AD risk SNPs (excluding *APOE* and *TOMM40* SNPs) per individual did not correlate with age of AD diagnosis ($p = 0.32$; **Figure 4E**). However, if netSNP is used to calculate *CVt* for each of these SNPs, the number of *CVt*-weighted SNPs did correlate inversely with age of AD diagnosis [$F(2,4750) = 419$, $\beta = -22$, $p < 2.2e-89$; **Figure 4E**], supporting the view that *CVt* provides a quantitative measure of the impact of an SNP on AD pathophysiology. We were concerned that the netSNP method may ascribe *CVt* values to SNPs based on genetic linkage to *APOE* $\epsilon 2$ or $\epsilon 4$, therefore we performed simulations using BD populations (see section “Materials and Methods”: *netSNP Validation Simulations*). These simulations support the view that the netSNP method is not choosing AD “protective” and AD “at-risk” SNPs based on genetic linkage or some other bias introduced in the netSNP procedure.

We also examined the relation of netSNP-identified *tSNPs* to the Braak scores that individuals (cases and controls) received during autopsy. The number of netSNP-identified “AD protective *tSNPs*” harbored per person displayed a negative correlation with Braak scores [$F(2,2698) = 349$, $\beta = 0.08$, $p < 3.0e-73$; **Figure 4H**], while the number of netSNP-identified “AD risk *tSNPs*” harbored per person displayed a positive correlation with Braak scores [$F(2,2698) = 272$, $\beta = 0.09$, $p < 3.4e-58$; **Figure 4H**]. These significant correlations, and the effect of providing

CVt weights, were obtained if *APOE* and *TOMM40* SNPs were (**Figure 4**) or were not (**Supplementary Figure 8**) included in the training matrix, indicating that the observed correlations were not driven by *APOE* (or SNPs in linkage disequilibrium with *APOE*; see **Supplementary Figure 13** and “Materials and Methods”).

DISCUSSION

Here we applied a standard and modified neural network tool to a large LOAD dataset and examined the association of SNPs to AD. We found that a standard NN trained with 50 SNPs can identify an individual’s cohort identity above chance; thus data were subsequently analyzed using only cohorts that were case:control balanced. Comparing Q-Q plots for AD and simulated constructed diseases (based on real genes with SNPs that have true population frequency as *APOE* $\epsilon 2$ and $\epsilon 4$) supports previous suggestions (Escott-Price et al., 2015, 2017) that there exist considerably more SNPs than the 20 previously identified as AD-associated.

An NN trained with 50 SNPs can predict dataset cases with accuracy greater (albeit, slightly) than if using only *APOE* SNPs genotypes, or a basic logistic regression model. NN accuracy was related approximately linearly with training set size, suggesting increasing dataset size will increase NN accuracy. NN accuracy was above chance if an NN was trained without (a) *APOE* SNPs, or (b) previously published AD-linked SNPs, or (c) 50 SNPs displaying the greatest case control asymmetry. These findings further support the view (Escott-Price et al., 2015, 2017) that more than the previously identified SNPs contain information regarding AD.

We developed netSNP, which investigated the impact of specific SNPs on NN output. In netSNP, once an NN was trained, the holdout set genotype was artificially assigned a single (or multiple) target SNP(s); in the general case the target SNP was assigned as homozygous to the minor frequency allele; the effect of the artificially introduced genotype was reflected by how much the NN output value was modified. netSNP recapitulated well the effect of different *APOE* genotypes on NN output. netSNP identified several hundred SNPs with weight values (i.e., *mCVt*) significantly outside values produced by randomly chosen SNPs. Some netSNP-identified SNPs had more extreme weight values than *APOE* $\epsilon 2$ or $\epsilon 4$. Notably, FishP values of SNPs with extreme *mCVt* values were not low in general, likely because too few individuals carry these SNPs. Yet their impact on NN output was large, possibly by leveraging non-linear interactions embedded in an NN. Notably, ADM (containing an SNP with the largest *mCVt* value despite a MAF = 0.009) was elevated in AD brains (Ferrero et al., 2017), contributed to age-related memory loss in mice (Larrayoz et al., 2017), was elevated in aging human brains (Larrayoz et al., 2017), and had been proposed as a novel drug target for AD (Ferrero et al., 2018). We also examined ABHD17A, as it relates to findings indicating that reduced function of this enzyme increases synaptic PSD-95 levels (Jeyifous et al., 2016; Yokoi et al., 2016), which protect

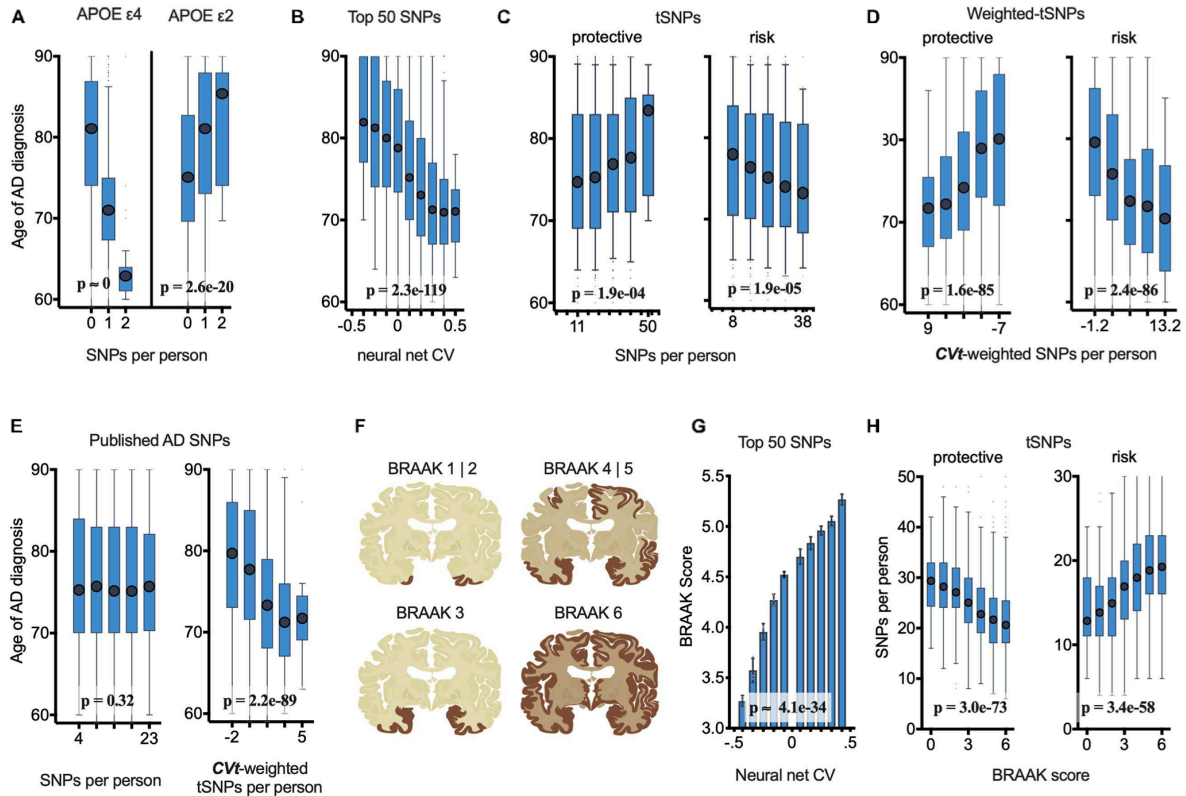


FIGURE 4 | netSNP validation: Number of netSNP-identified tSNPs and netSNP-CVt-weighted tSNPs correlate with age of AD diagnosis of all case individuals ($N = 4752$) and AD pathology (case and control receiving autopsy, $N = 2700$). (See Supplementary Figure 8 for results excluding APOE and TOMM40 in netSNP training matrices; results are similar; conclusions are the same.) (A) Age of AD diagnosis plotted against the number of APOE $\epsilon 4$ (or APOE $\epsilon 2$) SNPs per person. Here and below: boxplot X-axis values indicate mean value of boxed group; p -values based on general linear model analysis of variance. (B) Age of AD diagnosis plotted versus 50 SNP neural net CV. (C) Age of AD diagnosis plotted versus number of netSNP-identified AD-protective, left, or AD-risk (right) tSNPs. (D) Age of AD diagnosis plotted versus netSNP-CVt-weighted number of AD-protective (left) and AD-risk (right) tSNPs. (E) Age of AD diagnosis plotted versus number (left) or netSNP-CVt-weighted number (right) of previously published AD-linked SNPs, excluding those in APOE/TOMM40. n.s., not significant. (F) Diagram of human brain with affected regions for indicated Braak scores. (G) Neural net CV plotted versus Braak score. (H) netSNP-identified AD-protective (left) and AD-risk (right) tSNPs per person plotted versus Braak score.

synapses from beta amyloid (Malinow, *unpublished observation*). netSNP predicted that an *ABHD17A* SNP was protective for individuals with *APOE* $\epsilon 4$ (see **Supplementary Table 2**). Indeed, we found that $\epsilon 4$ carrier case individuals with this *ABHD17A* SNP received an AD diagnosis almost 6 years later than such individuals without this SNP [76.6 years ($N = 19$) vs. 70.8 years ($N = 1831$), $p < 0.0001$; t -test], which is consistent with this SNP being protective against AD in *APOE* $\epsilon 4$ carriers. These findings support the view that netSNP can identify AD-relevant SNPs.

To validate netSNP we considered variables not used in any netSNP computations: age of an individual's AD diagnosis (cf., Mars et al., 2020) and Braak score. The number of netSNP-identified "AD-protective SNPs" harbored by an individual correlated significantly with the age an individual was diagnosed with AD and inversely with Braak score; while the number of

netSNP-identified "AD-risk SNPs" harbored by an individual correlated significantly inversely with the age an individual was diagnosed with AD and positively with Braak score. Scaling each netSNP-identified SNP with CVt increased the significance of these correlations. Notably, applying netSNP-derived CVt weights to previously reported AD SNPs (each thought to have a small effect on AD pathophysiology) converted their correlation to age of diagnosis from not significant to significant, suggesting that netSNP can accurately assess small-effect SNPs. The correlations examined in this validation test hold if *APOE* or *TOMM40* are not used in the training step of netSNP, indicating that the netSNP-identified SNPs as well as the netSNP-generated CVt weights are not dependent on a bias imposed by *APOE* SNPs (or SNPs in linkage disequilibrium with *APOE*, **Supplementary Figure 13**) in netSNP. Further validation of netSNP and netSNP-identified SNPs suggested

to be “protective” or “at-risk” in this study will require tests using an independent AD dataset as well as biological experimentation.

Our data suggest the set, as a whole, of netSNP-identified SNPs are highly predictive of AD age of onset and physiological severity, and their relative importance may be indicated by the netSNP-derived *mCVt* weight. The netSNP-identified SNPs would each, on average, be expected to have a small impact on the disease (on average 1/200 that of *APOEε4*; but see above for *ABHD17A* SNP). Insight into AD provided by such small-effect SNPs will require computational methods that can analyze disease and biochemical pathways from large groups of genes. Such tools may be aided by incorporation of *mCVt* values.

In general, our findings suggest that netSNP may be useful in identifying pathophysiologically relevant genes in AD; it may be equally applicable to other conditions. It will be important to test these methods on a completely independent AD dataset with similar ethnic make-up (and compare those results with results in this study), as well as AD datasets with different ethnic backgrounds, for this method to be generally applicable to the multicultural nature of the United States and world population (Martin et al., 2019).

MATERIALS AND METHODS

Alzheimer's Disease Sequencing Project Dataset

The dataset used in these analyses was generously provided by the Alzheimer's Disease Sequencing Project (ADSP), and has been previously described in detail in other manuscripts (Harold et al., 2009; Raghavan et al., 2018) and online at niagads.org. To summarize, individuals in this dataset were from well-characterized cohorts, including 6,000 individuals diagnosed with late-onset Alzheimer's disease (mean age of diagnosis: 75.4) and 5,000 elderly controls without dementia (mean age: 86.1, at the date of last visit to AD practitioner). Whole-exome sequencing data for each individual went through a quality-control “cleaning” process by two independent sources (Baylor and Broad Institutes), and was provided in variant call format (.vcf); genotype data was accompanied by several phenotypic and qualitative metrics (e.g., each individual's sex, age, race, cohort, etc.). For 28% of individuals an autopsy was performed and their Braak staging score was reported (Braak et al., 2006). Data are available for download upon administrative approval from the NIA Genetics of Alzheimer's Disease Storage Site (NIAGADS).

VCF Data Compression

Raw SNP data were passed through an automated preprocessing pipeline that involved reducing the dataset size by 100-fold using sparse matrices and annotating SNPs of interest. The raw data were downloaded to a secure local hard drive as VCFs. VCFs were formatted as a matrix with rows being loci and columns being samples. This matrix was converted into a structure like an adjacency list. Sample IDs were replaced with seven-digit IDs.

Flags passed through the VCFs were converted to numeric flags. Counts of homozygous and heterozygous samples, as well as the sample names and genotypes were recorded per locus. The dataset was binned into three bins according to the following criteria: first, if the genotype was heterozygous (noted as 1), or homozygous (noted as 2) for the alternate allele. The second, if the genotype was homozygous for the reference allele (noted as 0). Third, if there was missing data for that sample (noted as 1). The combination of the bins and information contained within makes the 400-fold compression conversion a lossless process. The resulting matrices were relatively small and thus easier to query/manipulate than VCFs.

General Data Processing

Unless otherwise stated, data processing and analyses were conducted using MATLAB scientific computing software (Mathworks, 2020a,b). A compressed version of the data (as described in the section above) was imported into the MATLAB workspace. The data were then prepared for machine learning by splitting the data into training and holdout datasets. As the data were split, an attempt was made to balance cases and controls from each cohort. Cohorts that had too few cases or controls (<20% of each other; or fewer than 20 individuals) were omitted (see **Supplementary Figure 2**). After splitting and counterbalancing, a Fisher's exact test was performed for each SNP to assign a *p*-value to the case:control asymmetries. SNPs were then sorted, ascending, by *p*-value.

Artificial Neural Network Classification

In most instances, the model training matrix (feature matrix) consisted of individual genotypes for the 50 top SNPs after sorting SNPs by the training group's Fisher's exact test *p*-value. The rows and columns of this feature matrix represented individuals and SNPs, respectively, with each cell indicating whether a person was a homozygous reference, heterozygous, or homozygous alternate (see **Supplementary Figure 3**).

For polygenic classification we used a multilayer pattern recognition neural network (Mathworks, 2020a). This feed-forward neural net architecture can be trained to predict target classes (i.e., “labels” or “conditions” like case/control) based on a set of training features (Demuth et al., 2014). Labels for pattern recognition networks in a binary classification problem consist of a vector of 0s and 1s, where a 0 represents the negative condition (i.e., control), while a 1 represents the positive condition (i.e., case). In our formulation a pattern recognition network includes the following parameterization:

$$patternnet(nLayers, fTrain, fPerf)$$

where *nLayers* is the row vector of length *n*, representing the number of hidden layers; each *n*th value specifies the number of neurons in a given layer (e.g., 50, 10 would have two hidden layers of 50 neurons and 10 neurons, respectively). *fTrain* specifies the network training function (e.g., BFGS Quasi-Newton). *fPerf* specifies the performance function (e.g., cross-entropy).

We used a scaled conjugate gradient (SCG) training function for the polygenic classification task ($fTrain = SCG$). The SCG network training function updates network weights and bias values using conjugate gradient backpropagation, and can be used to train any network with derivatives for weight, input, and transfer functions (Moller, 1993). With regard to network training speed, SCG is significantly faster than other conjugate gradient methods, because it does not require line searches during each machine learning iteration (0.1 core hours per training session). Parameterization of the training function involves:

$fTrain$ ($maxEpochs$, $minGrad$, $maxFails$, $WtSigma$, $Lambda$) where $maxEpochs$ is the maximum number of epochs to train (e.g., 1000), $minGradient$ is the minimum performance gradient (e.g., $1e-6$), $maxFails$ is the maximum validation failures allowed (e.g., 10), $WtSigma$ is the change in weight for second derivative approximation (e.g., $5.0e-5$), and $Lambda$ regulates the indefiniteness of the Hessian (e.g., $5.0e-7$). Unless otherwise noted, the model was implemented in the MATLAB (*Mathworks – Deep Learning Toolbox*) scientific programming environment and parameterized with the following values:

```
patternnet(nLayers = (50, 10), fTrain = "SCG",
  fPerf = "cross-entropy")
SCG(maxEpochs = 1000, minGrad = 1e-6, maxFails = 10,
  WtSigma = 5e-5, Lambda = 5e-7)
cross-entropy(reg = 0.1, norm = (-0.5, 0.5)).
```

The last steps involve preparing the data for network training: (1) individuals are randomly split into a training, validation, or holdout group; (2) a Fisher's exact test is used to compute the p -value associated with the case:control asymmetry in the training set at each variant locus; (3) the list of SNPs are sorted, ascending by p -value; and (4), some number of SNPs (e.g., the top 50) are selected for generating an individual-by-SNP matrix, where each cell contains the genotype of a given person at a given SNP locus. Finally, with the feature matrices prepared, and the model fully parameterized, neural net training can commence:

```
net = train(patternnet, Xt, Yt, Xv, Yv)
```

Again, *patternnet* represents the parameterized model (and all instructions for model training), Xt and Xv represent the individual-by-SNP feature matrix for the training and validation groups, respectively, and Yt and Yv are binary arrays indicating whether each person is a case or control (i.e., the condition labels). The model is trained as described above, and the final output is a fitted neural network model (a set of network weights).

netSNP Validation Test Using BD Populations

We conducted simulations to rule out the possibility that the netSNP method may choose SNPs based on genetic linkage to APOE $\epsilon 2$ or $\epsilon 4$; i.e., significant tSNPs could

display at-risk or protective properties despite their not being pathophysiologically associated with AD. Furthermore, other details of the netSNP method may predispose cases to artifactual correlations with age of AD diagnosis and Braak scores (We note, however, that neither the age of AD diagnosis, nor their Braak score, was used in any calculations performed in section "NN Construction and Performance" or "netSNP Description and Application").

We thus tested for the correlations shown in section "NN and CV as Predictors of AD and Its Pathophysiology." for BDs 1–12 (see above; **Supplementary Table 3** and **Supplementary Figure 7**). Age of diagnosis of BD was ascribed based on APOE SNPs effects in age of AD diagnosis (using MATLAB empirical cumulative distribution functions). For each BD, a balanced dataset was constructed (as for AD, see section "Dataset Pipeline, Case:Control Balancing and SNP Properties"), and BD "protective" and "at-risk" tSNPs were identified as described for AD in section "NN and CV as Predictors of AD and Its Pathophysiology." Next, we considered the set of individuals ascribed BD. We computed a correlation probability, based on a general linear model, between their age of BD diagnosis and the number of BD tSNPs or number of BDCVt-weighted tSNPs. Results for one BD (based on a BD constructed from BDgene CHSY1; **Supplementary Figure 7D**) is compared with results for AD (**Supplementary Figure 7C**). A summary of results for the 12 separate BDs, and AD for comparison, are shown in **Supplementary Table 3**. Note that for no BD was there a significant correlation (right columns). These simulations support the view that the netSNP method is not choosing AD "protective" and AD "at-risk" SNPs based on genetic linkage or some other bias introduced in the netSNP procedure.

Statistics

Statistical methods described per figure below.

For each BD constructed, individuals in the ADSP population were assigned a BD based on their genotype; those with APOE $\epsilon 2$ -like SNPs were randomly assigned as control with OR 2.41; those with APOE $\epsilon 4$ -like SNPs were assigned as case with OR 0.30. Those without either SNPs were assigned randomly to control with OR 0.89 (see **Supplementary Table 3**). To generate random Q-Q plots, 100 datasets were generated with randomly scrambled case-control labels. Fisher's exact test p -values were then computed for those 100 scrambled sets. Scrambled sets were plotted against each other to generate the C.I. region (gray dots) and also plotted against the actual data (colored dots).

Hundred random groups were generated with cases and controls counterbalanced within cohorts to formulate neural network training matrices. As described above, in each run one of these random groups was selected and an artificial NN was trained using the 50 SNPs with the lowest Fisher's exact test p -value among training group individuals. NN classifier performance on the holdout set was then evaluated. A histogram of each individual's mean NN classifier value (CV). Shows receiver operator characteristic (ROC) curves using SNP sets as features and normalizing CVs to range between 0 and 1: curve

“1–50” used 50 SNPs with the lowest training group p -values; “1–50 -*APOE*” used 50 SNPs with the lowest training group p -values omitting *APOE* and *TOMM40*; “1–50 -GWAS” used 50 SNPs with the lowest training group p -values omitting SNPs that previously met genome-wide significance in the literature; “51–100” used SNPs with the 51st–100th lowest training group p -values. The left panel shows the mean correct predictions in percent for each condition in (Figure 2B); the right panel was generated like the left panel’s “Top 50,” except the experimental manipulation varied the number of samples in the training group (1 sample = 1 chromosome), as indicated in the figure legend. CVs were generated like in “2B 1–50,” and normalized to a range between –0.5 and 0.5. The classification threshold was fixed at 0 and the false discovery rate (FDR) and positive predictive value (PPV) were then computed at each corresponding x -axis case prevalence. The FDR and PPV were also computed using the optimal operating point (OOP):

$$S = \frac{\text{Cost}(P|N) - \text{Cost}(N|N)}{\text{Cost}(N|P) - \text{Cost}(P|P)} * \frac{N}{P}$$

where $\text{Cost}(N|P)$ is the cost of misclassifying a case, $\text{Cost}(P|N)$ is the cost of misclassifying a control, where $P = TP + FN$, and $N = TN + FP$ (TP , true positive; TN , true negative; FP , false positive; FN , false negative). The OOP was then determined by moving a line with slope S from $FPR = 0$, $TPR = 1$ (the top left of the ROC) down-and-right, until it intersected with the ROC curve (Mathworks, 2020b).

The histograms shown in (Figure 3A) are the result of training an NN using individuals of all *APOE* subtypes, and applying this NN on holdout set individuals assigned to each of the six *APOE* genotypes. That is, after the NN is trained as described above in *General Data Preprocessing*, all holdout individuals are assigned the *APOE*ε22 genotype and a histogram is generated; then all holdout individuals are assigned the *APOE*ε23 genotype and another histogram is generated, etc. We call this genotype assignment procedure the *netSNP method* (described below) which we show can be used as a general method for assessing the importance of any SNP on NN performance. For comparison, histograms shown in (Figure 3B) are the result of training an NN using a balanced set of individuals, and computing CVs for holdout set subgroups of individuals with the *APOE* genotypes limited to one of *APOE*ε22, ε23, ε24, ε33, ε34, or ε44. **netSNP method:** 4,000 target SNPs were chosen based on them having the lowest Fisher’s exact test p -value. For each target SNP, the netSNP method can produce a *NAT*, *REF*, *ALT*, and *DIF* value for each individual. For a single target SNP, obtaining these values involved the following steps: (1) a target SNP was selected to be part of a 50-SNP training matrix. (2) A random subset (~70%) of a balanced set of individuals served as a training set. (3) From this training set, a Fisher’s exact test p -value was calculated for each of the (1.4 million) SNPs. (4) A single target SNP was paired with the 49 SNPs with the lowest p -value to generate a neural network training matrix. (5) The neural network was trained as described above in the *General Data Processing* methods. (6) A CV score was

generated for each of the individuals in the holdout set (*NAT* score). (7) All holdout individuals were assigned the homozygous reference genotype for the target SNP and again a CV was generated (*REF* score). (8) All holdout individuals were assigned the homozygous alternate allele (minor frequency allele) for the target SNP and a CV was generated (*ALT* score). (9) The difference between the *ALT* and *REF* scores were computed (*DIF* score). This procedure was performed 20x for each target SNP; for a given target SNP, each individual’s average *ALT* score represents that individual’s **CV_t** score for the given target SNP. In this study we tested if **CV_t** value could be considered a weighted measure of the impact of target SNP t on the NN. Similar to how histograms are generated for, after the NN was trained as described above in *General Data Preprocessing*, all holdout individuals were assigned the homozygous genotype for minor frequency allele of the target SNP for the indicated gene (see Table 1 for chromosome and position of the target SNP for each indicated gene).

Boxplots in (Figure 4A) were generated by grouping case individuals based on whether they had a homozygous reference, heterozygous, or homozygous minor frequency for the indicated allele, and plotted the median AD age-of-onset (+/- interquartile range, IQR; whiskers = range; dots = outliers). Boxplots were generated by pooling case individuals into six bins that were uniformly discretized based on the NN CV value, on the number of protective (left) or risk (right) target SNPs each individual had, or **CV_t**-weighted target SNPs, and then plotted the median AD age-of-onset (+/- IQR; whiskers = range; dots = outliers) for each of these bins. Figure 4E (left) was generated like Figure 4C, considering previously published (without *APOE*) AD SNPs. Figure 4D (right) was generated like Figure 4C, providing a netSNP-computed **CV_t** for each previously published (without *APOE*) AD SNPs. Brain sections in (Figure 4F) depict Braak staging – a method used to classify the degree of pathology in Alzheimer’s disease – commonly used in post-mortem clinical diagnosis of AD by performing brain autopsy; images here intend to summarize the general disease sequelae as shown in actual brain images from Braak et al. (2006). The bar plot in (Figure 4G) was generated by identifying individuals that had **mCV_t** scores across all **tSNPs** that fall into each of the indicated bins, and the mean Braak stage of the individuals in each bin was plotted. Boxplots in (Figure 4H) pool individuals based on ADSP-reported Braak values and plot the median number of target SNPs (+/- IQR; whiskers = range; dots = outliers) found in individuals with a brain pathology that fall into one of these six Braak stages; as in (Figures 4C,D), effects are shown separately for SNPs that potentially confer protection (left panel) and risk (right panel). P -values were computed using a general linear model, where p -value represents the probability of the slope coefficient having such a magnitude, under the null hypothesis.

DATA AVAILABILITY STATEMENT

The data analyzed in this study is subject to the following licenses/restrictions: NIH ADSP Embargo – Access granted via application. Requests to access these datasets should be

directed to <https://dss.niagads.org/documentation/applying-for-data/application-instructions/>.

ETHICS STATEMENT

Ethical review and approval was not required for the study on human participants in accordance with the local legislation and institutional requirements. The patients/participants provided their written informed consent to participate in this study.

AUTHOR CONTRIBUTIONS

BM and RM designed the study. BM, SP, TG, and RM prepared and preprocessed the data. BM, AnR, AIR, and RM performed the statistical analyses. BM and RM generated figures. BM, AIR, and RM wrote the manuscript. All authors contributed to the article and approved the submitted version.

FUNDING

D. H. Chen Foundation Grant “Preventing Alzheimer’s Disease,” NIH1R01-EY022306. The Alzheimer’s Disease Sequencing Project (ADSP) is comprised of two Alzheimer’s Disease (AD) genetics consortia and three National Human Genome Research Institute (NHGRI) funded Large Scale Sequencing and Analysis Centers (LSAC). The two AD genetics consortia are the Alzheimer’s Disease Genetics Consortium (ADGC) funded by NIA (U01 AG032984), and the Cohorts for Heart and Aging Research in Genomic Epidemiology (CHARGE) funded by NIA (R01 AG033193), the National Heart, Lung, and Blood Institute (NHLBI), other National Institutes of Health (NIH) institutes, and other foreign governmental and non-governmental organizations. The Discovery Phase analysis of sequence data is supported through UF1AG047133 (to Schellenberg, Farrer, Pericak-Vance, Mayeux, and Haines); U01AG049505 to Seshadri; U01AG049506 to Boerwinkle; U01AG049507 to Wijsman; and U01AG049508 to Goate, and the Discovery Extension Phase analysis is supported through U01AG052411 to Goate, U01AG052410 to Pericak-Vance, and U01AG052409 to Seshadri and Fornage. Data generation and harmonization in the Follow-up Phases is supported by U54AG052427 (to Schellenberg and Wang). The ADGC cohorts include: Adult Changes in Thought (ACT), the Alzheimer’s Disease Centers (ADCs), the Chicago Health and Aging Project (CHAP), the Memory and Aging Project (MAP), Mayo Clinic (MAYO), Mayo Parkinson’s Disease controls, University of Miami, the Multi-Institutional Research in Alzheimer’s Genetic Epidemiology Study (MIRAGE), the National Cell Repository for Alzheimer’s Disease (NCRAD), the National Institute on Aging Late Onset Alzheimer’s Disease Family Study (NIA-LOAD), the Religious Orders Study (ROS), the Texas Alzheimer’s Research and Care Consortium (TARCC), Vanderbilt University/Case Western Reserve University (VAN/CWRU), the Washington Heights-Inwood Columbia Aging Project (WHICAP) and the Washington University

Sequencing Project (WUSP), the Columbia University Hispanic-Estudio Familiar de Influencia Genetica de Alzheimer (EFIGA), the University of Toronto (UT), and Genetic Differences (GD). The CHARGE cohorts are supported in part by National Heart, Lung, and Blood Institute (NHLBI) infrastructure grant HL105756 (Psaty), RC2HL102419 (Boerwinkle), and the neurology working group is supported by the National Institute on Aging (NIA) R01 grant AG033193. The CHARGE cohorts participating in the ADSP include the following: Austrian Stroke Prevention Study (ASPS), ASPS-Family study, and the Prospective Dementia Registry-Austria (ASPS/PRODEM-Aus), the Atherosclerosis Risk in Communities (ARIC) Study, the Cardiovascular Health Study (CHS), the Erasmus Rucphen Family Study (ERF), the Framingham Heart Study (FHS), and the Rotterdam Study (RS). ASPS is funded by the Austrian Science Fund (FWF) grant numbers P20545-P05 and P13180 and the Medical University of Graz. The ASPS-Fam is funded by the Austrian Science Fund (FWF) project I904, the EU Joint Programme – Neurodegenerative Disease Research (JPND) in frame of the BRIDGET project (Austria, Ministry of Science) and the Medical University of Graz and the Steiermärkische Krankenanstaltengesellschaft. PRODEM-Austria is supported by the Austrian Research Promotion Agency (FFG) (Project No. 827462) and by the Austrian National Bank (Anniversary Fund, project 15435). ARIC research was carried out as a collaborative study supported by NHLBI contracts (HHSN268201100005C, HHSN268201100006C, HHSN268201100007C, HHSN268201100008C, HHSN268201100009C, HHSN268201100010C, HHSN268201100011C, and HHSN268201100012C). Neurocognitive data in ARIC are collected by U01 2U01HL096812, 2U01HL096814, 2U01HL096899, 2U01HL096902, 2U01HL096917 from the NIH (NHLBI, NINDS, NIA, and NIDCD), and with previous brain MRI examinations funded by R01-HL70825 from the NHLBI. CHS research was supported by contracts HHSN268201200036C, HHSN268200800007C, N01HC55222, N01HC85079, N01HC85080, N01HC85081, N01HC85082, N01HC85083, N01HC85086, and grants U01HL080295 and U01HL130114 from the NHLBI with additional contribution from the National Institute of Neurological Disorders and Stroke (NINDS). Additional support was provided by R01AG023629, R01AG15928, and R01AG20098 from the NIA. FHS research was supported by NHLBI contracts N01-HC-25195 and HHSN268201500001I. This study was also supported by additional grants from the NIA [R01s AG054076, AG049607, and AG033040, and NINDS (R01 NS017950)]. The ERF study as a part of EUROSPAN (European Special Populations Research Network) was supported by European Commission FP6 STRP grant number 018947 (LSHG-CT-2006-01947) and also received funding from the European Community’s Seventh Framework Programme (FP7/2007-2013)/grant agreement HEALTH-F4-2007-201413 by the European Commission under the programme “Quality of Life and Management of the Living Resources” of 5th Framework Programme (no. QL2-CT-2002-01254). High-throughput analysis of the ERF data was supported by a joint grant from the Netherlands Organisation for Scientific Research and the Russian Foundation

for Basic Research (NWO-RFBR 047.017.043). The Rotterdam Study is funded by Erasmus University Medical Center and Erasmus University Rotterdam, the Netherlands Organisation for Health Research and Development (ZonMw), the Research Institute for Diseases in the Elderly (RIDE), the Ministry of Education, Culture and Science, the Ministry for Health, Welfare and Sports, the European Commission (DG XII), and the municipality of Rotterdam. Genetic datasets are also supported by the Netherlands Organisation for Scientific Research NWO Investments (175.010.2005.011, 911-03-012), the Genetic Laboratory of the Department of Internal Medicine, Erasmus MC, the Research Institute for Diseases in the Elderly (014-93-015; RIDE2), and the Netherlands Genomics Initiative (NGI)/Netherlands Organisation for Scientific Research (NWO), Netherlands Consortium for Healthy Aging (NCHA), project 050-060-810. All studies are grateful to their individuals, faculty, and staff. The content of these manuscripts is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health or the U.S. Department of Health and Human Services. The four LSACs are: the Human Genome Sequencing Center at the Baylor College of Medicine (U54 HG003273), the Broad Institute Genome Center (U54HG003067), The American Genome Center at the Uniformed Services University of the Health Sciences (U01AG057659), and the Washington University Genome Institute (U54HG003079). Biological samples and associated phenotypic data used in primary data analyses were stored at Study Investigators institutions, and at the National Cell Repository for Alzheimer's Disease (NCRAD, U24AG021886) at Indiana University funded by NIA. Associated Phenotypic Data used in primary and secondary data analyses were provided by Study Investigators, the NIA funded Alzheimer's Disease Centers (ADCs), and the National Alzheimer's Coordinating Center (NACC, U01AG016976) and the National Institute on Aging Genetics of Alzheimer's Disease Data Storage Site (NIAGADS, U24AG041689) at the University of Pennsylvania, funded by NIA, and at the Database for Genotypes and Phenotypes (dbGaP) funded by NIH. This research was supported in part by the Intramural Research Program of the National Institutes of Health, National Library of Medicine. Contributors to the Genetic Analysis Data included Study Investigators on projects that were individually funded by NIA, and other NIH institutes, and by private United States organizations, or foreign governmental or non-governmental organizations.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fgene.2021.647436/full#supplementary-material>

REFERENCES

- Altmann, A., Tolo, S. I., Sander, O., and Lengauer, T. (2010). Permutation importance: a corrected feature importance measure. *Bioinformatics* 26, 1340–1347. doi: 10.1093/bioinformatics/btq134
- Beecham, G. W., Bis, J. C., Martin, E. R., Choi, S.-H., DeStefano, A. L., van Duijn, C. M., et al. (2017). The Alzheimer's disease sequencing project:

Supplementary Figure 1 | A multi-layer feed forward neural network was trained to classify individuals by cohort identity. The scaled conjugate gradient (SCG) algorithm was the primary learning algorithm used to minimize neural network weights, here and in all other applications of a neural network in this manuscript unless otherwise noted.

Supplementary Figure 2 | ROC curves of NN performance were generated based on neural net output after being trained on individual cohort labels, and using individual genotypes (for the top 50 SNPs) as training features.

Supplementary Figure 3 | For neural net architecture used to classify individuals as cases or controls see methods for Supplementary Figure 1.

Supplementary Figure 4 | See legend.

Supplementary Figure 5 | Distributions of CVI values are based on netSNP outputs (see netSNP method above). To determine if any tSNPs might confer protection from APOE4 effects, neural nets weights were fit using true training group genotypes; then the netSNP test was performed simultaneously for APOE4 and the tSNP, such that in each holdout group individual the APOE locus was set to $\epsilon 4$ and the target SNP locus were set to the homozygous minor alleles, and neural net output (CVI) was evaluated.

Supplementary Figure 6 | The CVI distributions of holdout group individuals for different tSNPs were generated using the netSNP method (see netSNP method above); here APOE and TOMM40 were excluded as training features.

Supplementary Figure 7 | Same methods used to generate Figure 4, with the only modification being that APOE and TOMM40 were excluded as training features, except Panel-A which necessarily includes APOE as a training feature. The left boxplot of Panel-E is the same as in Figure 4 since the raw per-person count of known AD genes is independent of neural network output and netSNP manipulations.

Supplementary Figure 8 | See text: results section "netSNP as Predictor of AD-Linked tSNPs and AD Pathophysiology." For cases, ages of BD diagnosis was assigned based on AD age distribution for analogous genotype [e.g., those with SNP2 (E4-like) were given age of diagnosis with a distribution as individuals with APOE4 are diagnosed with AD]. For each BD, netSNP analysis was performed as it was for AD (netSNP method, above). Correlations between age of BD diagnosis and 12 of netSNP-identified SNPs and CVI-weighted 12 of netSNP-identified SNPs was conducted as for AD (Figure 4).

Supplementary Figure 9 | Same methods used to generate Figure 2D using a classification threshold fixed at zero. NCASE, the number of individuals in the case condition; PCASE, the number of individuals predicted as case; NCTRL, the number of individuals in the control condition; PCTRL, the number of individuals predicted as control; TP, true positive; TN, true negative; FP, false positive; FN, false negative; TPR, true positive rate; TNR, true negative rate; FPR, false positive rate; FNR, false negative rate; PPV, positive predictive value; NPV, negative predictive value; FOR, false omission rate; FDR, false discovery rate; AUC, area under the curve; ACC, accuracy; FOS, F1-score; PLR, positive likelihood ratio; NLR, negative likelihood ratio; DOR, diagnostic odds ratio; MCC, Matthews correlation coefficient.

Supplementary Figure 10 | Same methods used to generate Figure 2D using the OOP as the classification threshold. Same abbreviations as in Supplementary Figure 9.

Supplementary Figure 11 | Definitions in Supplementary Figure 9 shown in the form of a confusion matrix.

Supplementary Figure 12 | See Figure 1C.

Supplementary Figure 13 | See Figure legend.

study design and sample selection. *Neurol. Genet.* 3:e194. doi: 10.1212/nxg.000000000000194

- Beecham, G. W., Vardarajan, B., Blue, E., Bush, W., Jaworski, J., Barral, S., et al. (2018). Rare genetic variation implicated in non-Hispanic white families with Alzheimer disease. *Neurol. Genet.* 4:e286. doi: 10.1212/nxg.0000000000000286
- Bis, J. C., Jian, X., Kunkle, B. W., Chen, Y., Hamilton-Nelson, K. L., Bush, W. S., et al. (2018). Whole exome sequencing study identifies novel rare and

- common Alzheimer's-associated variants involved in immune response and transcriptional regulation. *Mol. Psychiatr.* 25, 1–17. doi: 10.1038/s41380-018-0112-7
- Braak, H., Alafuzoff, I., Arzberger, T., Kretschmar, H., and Tredici, K. D. (2006). Staging of Alzheimer disease-associated neurofibrillary pathology using paraffin sections and immunocytochemistry. *Acta neuropathol.* 112, 389–404. doi: 10.1007/s00401-006-0127-z
- Bulik-Sullivan, B. K., Loh, P.-R., Finucane, H. K., Ripke, S., Yang, J., Patterson, N., et al. (2015). LD Score regression distinguishes confounding from polygenicity in genome-wide association studies. *Nat. Genet.* 47, 291–295. doi: 10.1038/ng.3211
- Carmona, S., Hardy, J., and Guerreiro, R. (2018). The genetic landscape of Alzheimer disease. *Handb. Clin. Neurol.* 148, 395–408. doi: 10.1016/b978-0-444-64076-5.00026-0
- Cauwenbergh, C. V., Broeckhoven, C. V., and Sleegers, K. (2015). The genetic landscape of Alzheimer disease: clinical implications and perspectives. *Genet. Med.* 18, 421–430. doi: 10.1038/gim.2015.117
- Corder, E., Saunders, A., Strittmatter, W., Schmechel, D., Gaskell, P., Small, G., et al. (1993). Gene dose of apolipoprotein E type 4 allele and the risk of Alzheimer's disease in late onset families. *Science* 261, 921–923. doi: 10.1126/science.8346443
- Crane, P. K., Foroud, T., Montine, T. J., and Larson, E. B. (2017). Alzheimer's disease sequencing project discovery and replication criteria for cases and controls: data from a community-based prospective cohort study with autopsy follow-up. *Alzheimers Dement.* 13, 1410–1413. doi: 10.1016/j.jalz.2017.09.010
- Demuth, H. B., Beale, M. H., Jess, O. D., and Hagan, M. T. (2014). *Neural Network Design*. Available online at: <https://dl.acm.org/doi/book/10.5555/2721661> (accessed February 24, 2021).
- Desikan, R. S., Fan, C. C., Wang, Y., Schork, A. J., Cabral, H. J., Cupples, L. A., et al. (2017). Genetic assessment of age-associated Alzheimer disease risk: development and validation of a polygenic hazard score. *PLoS Med.* 14:e1002258. doi: 10.1371/journal.pmed.1002258
- DeTure, M. A., and Dickson, D. W. (2019). The neuropathological diagnosis of Alzheimer's disease. *Mol. Neurodegener.* 14:32. doi: 10.1186/s13024-019-0333-5
- Escott-Price, V., Myers, A. J., Huentelman, M., and Hardy, J. (2017). Polygenic risk score analysis of pathologically confirmed Alzheimer disease. *Ann. Neurol.* 82, 311–314. doi: 10.1002/ana.24999
- Escott-Price, V., Sims, R., Bannister, C., Harold, D., Vronskaya, M., Majounie, E., et al. (2015). Common polygenic variation enhances risk prediction for Alzheimer's disease. *Brain* 138, 3673–3684. doi: 10.1093/brain/awv268
- Ferrero, H., Larrayoz, I. M., Gil-Bea, F. J., Martínez, A., and Ramírez, M. J. (2018). Adrenomedullin, a novel target for neurodegenerative diseases. *Mol. Neurobiol.* 55, 8799–8814. doi: 10.1007/s12035-018-1031-y
- Ferrero, H., Larrayoz, I. M., Martisova, E., Solas, M., Howlett, D. R., Francis, P. T., et al. (2017). Increased levels of brain adrenomedullin in the neuropathology of Alzheimer's disease. *Mol. Neurobiol.* 55, 5177–5183. doi: 10.1007/s12035-017-0700-6
- Gatz, M., Reynolds, C. A., Fratiglioni, L., Johansson, B., Mortimer, J. A., Berg, S., et al. (2006). Role of genes and environments for explaining Alzheimer disease. *Arch. Gen. Psychiatr.* 63:168. doi: 10.1001/archpsyc.63.2.168
- Goldstein, B. A., Polley, E. C., and Briggs, F. B. S. (2011). Random forests for genetic association studies. *Stat. Appl. Genet. Mol.* 10:32. doi: 10.2202/1544-6115.1691
- Guerreiro, R. J., and Hardy, J. (2012). TOMM40 association with Alzheimer disease: tales of APOE and linkage disequilibrium. *Arch. Neurol.* 69, 1243–1244. doi: 10.1001/archneurol.2012.1935
- Gustaw-Rothenberg, K., Lerner, A., Bonda, D. J., Lee, H., Zhu, X., Perry, G., et al. (2010). Biomarkers in Alzheimer's disease: past, present and future. *Biomark Med.* 4, 15–26. doi: 10.2217/bmm.09.86
- Hampel, H., O'Bryant, S. E., Molinuevo, J. L., Zetterberg, H., Masters, C. L., Lista, S., et al. (2018). Blood-based biomarkers for Alzheimer disease: mapping the road to the clinic. *Nat. Rev. Neurol.* 14, 639–652. doi: 10.1038/s41582-018-0079-7
- Harold, D., Abraham, R., Hollingworth, P., Sims, R., Gerrish, A., Hamshere, M. L., et al. (2009). Genome-wide association study identifies variants at CLU and PICALM associated with Alzheimer's disease. *Nat. Genet.* 41, 1088–1093. doi: 10.1038/ng.440
- Jack, C. R., Bennett, D. A., Blennow, K., Carrillo, M. C., Dunn, B., Haeberlein, S. B., et al. (2018). NIA-AA research framework: toward a biological definition of Alzheimer's disease. *Alzheimers Dement.* 14, 535–562. doi: 10.1016/j.jalz.2018.02.018
- Jansen, I. E., Savage, J. E., Watanabe, K., Bryois, J., Williams, D. M., Steinberg, S., et al. (2019). Genome-wide meta-analysis identifies new loci and functional pathways influencing Alzheimer's disease risk. *Nat. Genet.* 51, 404–413. doi: 10.1038/s41588-018-0311-9
- Jeyifous, O., Lin, E. I., Chen, X., Antinone, S. E., Mastro, R., Drisdel, R., et al. (2016). Palmitoylation regulates glutamate receptor distributions in postsynaptic densities through control of PSD95 conformation and orientation. *Proc. Natl. Acad. Sci. U.S.A.* 113, E8482–E8491. doi: 10.1073/pnas.1612963113
- Jolliffe, I. T. (1986). *Principal Component Analysis*. Springer Series in Statistics. New York, NY: Springer, 115–128.
- Karch, C. M., and Goate, A. M. (2014). Alzheimer's disease risk genes and mechanisms of disease pathogenesis. *Biol. Psychiatr.* 77, 43–51. doi: 10.1016/j.biopsych.2014.05.006
- Koen, J. D., Barrett, F. S., Harlow, I. M., and Yonelinas, A. P. (2016). The ROC toolbox: a toolbox for analyzing receiver-operating characteristics derived from confidence ratings. *Behav. Res. Methods* 49, 1399–1406. doi: 10.3758/s13428-016-0796-z
- Koffie, R. M., Hashimoto, T., Tai, H.-C., Kay, K. R., Serrano-Pozo, A., Joyner, D., et al. (2012). Apolipoprotein E4 effects in Alzheimer's disease are mediated by synaptotoxic oligomeric amyloid- β . *Brain J. Neurol.* 135, 2155–2168. doi: 10.1093/brain/awv127
- Kunkle, B. W., Grenier-Boley, B., Sims, R., Bis, J. C., Damotte, V., Naj, A. C., et al. (2019). Genetic meta-analysis of diagnosed Alzheimer's disease identifies new risk loci and implicates A β , tau, immunity and lipid processing. *Nat. Genet.* 51, 414–430. doi: 10.1038/s41588-019-0358-2
- Lambert, J.-C., Heath, S., Even, G., Campion, D., Sleegers, K., Hiltunen, M., et al. (2009). Genome-wide association study identifies variants at CLU and CR1 associated with Alzheimer's disease. *Nat. Genet.* 41, 1094–1099. doi: 10.1038/ng.439
- Lander, E., and Schork, N. (1994). Genetic dissection of complex traits. *Science* 265, 2037–2048. doi: 10.1126/science.8091226
- Larrayoz, I. M., Ferrero, H., Martisova, E., Gil-Bea, F. J., Ramirez, M. J., and Martínez, A. (2017). Adrenomedullin contributes to age-related memory loss in mice and is elevated in aging human brains. *Front. Mol. Neurosci.* 10:384. doi: 10.3389/fnmol.2017.00384
- Ma, Y., Jun, G. R., Zhang, X., Chung, J., Naj, A. C., Chen, Y., et al. (2019). Analysis of Whole-exome sequencing data for Alzheimer disease stratified by APOE genotype. *JAMA Neurol.* 76, 1099–1108. doi: 10.1001/jamaneurol.2019.1456
- Mars, N., Koskela, J. T., Ripatti, P., Kiiskinen, T. T. J., Havulinna, A. S., Lindbohm, J. V., et al. (2020). Polygenic and clinical risk scores and their impact on age at onset and prediction of cardiometabolic diseases and common cancers. *Nat. Med.* 26, 549–557. doi: 10.1038/s41591-020-0800-0
- Martin, A. R., Kanai, M., Kamatani, Y., Okada, Y., Neale, B. M., and Daly, M. J. (2019). Clinical use of current polygenic risk scores may exacerbate health disparities. *Nat. Genet.* 51, 584–591. doi: 10.1038/s41588-019-0379-x
- Mathworks (2020a). *Multilayer Shallow Neural Network Architecture*. Mathworks: Natick, MA.
- Mathworks (2020b). *Receiver Operating Characteristic (ROC) Curve or Other Performance Curve for Classifier Output*. Mathworks: Natick, MA.
- Mendez, M. F. (2017). Early-Onset Alzheimer disease. *Neurol. Clin.* 35, 263–281. doi: 10.1016/j.ncl.2017.01.005
- Moller, M. F. (1993). A scaled conjugate-gradient algorithm for fast supervised learning. *Neural Netw.* 6, 525–533. doi: 10.1016/s0893-6080(05)80056-5
- Molnar, C. (2021). *Interpretable Machine Learning*. Available online at: <https://www.lulu.com/shop/christoph-molnar/interpretable-machine-learning/paperback/product-24036234.html> (accessed February 24, 2021).
- Naj, A. C., Lin, H., Vardarajan, B. N., White, S., Lancour, D., Ma, Y., et al. (2018). Quality control and integration of genotypes from two calling pipelines for whole genome sequence data in the Alzheimer's disease sequencing project. *Genomics* 111, 808–818. doi: 10.1016/j.ygeno.2018.05.004
- Patel, D., Mez, J., Vardarajan, B. N., Staley, L., Chung, J., Zhang, X., et al. (2019). Association of rare coding mutations with Alzheimer disease and other dementias among adults of European ancestry. *JAMA Netw. Open* 2:e191350. doi: 10.1001/jamanetworkopen.2019.1350
- Penner, G., Lecocq, S., Chopin, A., Vedoya, X., Lista, S., Vergallo, A., et al. (2019). Blood-based diagnostics of Alzheimer's disease. *Expert Rev. Mol. Diagn.* 19, 613–621. doi: 10.1080/14737159.2019.1626719
- Raghavan, N. S., Brickman, A. M., Andrews, H., Manly, J. J., Schupf, N., Lantigua, R., et al. (2018). Whole-exome sequencing in 20,197 persons for rare variants in Alzheimer's disease. *Ann. Clin. Trans. Neurol.* 5, 832–842.

- Saunders, A. M., Strittmatter, W. J., Schmechel, D., George-Hyslop, P. H. S., Pericak-Vance, M. A., Joo, S. H., et al. (1993). Association of apolipoprotein E allele 4 with late-onset familial and sporadic Alzheimer's disease. *Neurology* 43, 1467–1472. doi: 10.1212/wnl.43.8.1467
- Sejnowski, T. J. (2020). The unreasonable effectiveness of deep learning in artificial intelligence. *Proc. Natl. Acad. Sci. U.S.A.* 117:201907373. doi: 10.1073/pnas.1907373117
- Selzam, S., Coleman, J. R. I., Caspi, A., Moffitt, T. E., and Plomin, R. (2018). A polygenic p factor for major psychiatric disorders. *Transl. Psychiat.* 8:205. doi: 10.1038/s41398-018-0217-4
- Shi, Y., Yamada, K., Liddelow, S. A., Smith, S. T., Zhao, L., Luo, W., et al. (2017). ApoE4 markedly exacerbates tau-mediated neurodegeneration in a mouse model of tauopathy. *Nature* 549, 523–527. doi: 10.1038/nature24016
- Sisodia, S. S., and George-Hyslop, P. H. S. (2002). γ -Secretase, notch, A β and alzheimer's disease: where do the presenilins fit in? *Nat. Rev. Neurosci.* 3, 281–290. doi: 10.1038/nrn785
- Slatkin, M. (2007). Inbreeding coefficients and coalescence times. *Genet. Res.* 89, 479–487. doi: 10.1017/s0016672308009671
- Yokoi, N., Fukata, Y., Sekiya, A., Murakami, T., Kobayashi, K., and Fukata, M. (2016). Identification of PSD-95 dephalmitoylating enzymes. *J. Neurosci.* 36, 6431–6444. doi: 10.1523/jneurosci.0419-16.2016
- Yu, C.-E., Seltman, H., Peskind, E. R., Galloway, N., Zhou, P. X., Rosenthal, E., et al. (2007). Comprehensive analysis of APOE and selected proximate markers for late-onset Alzheimer's disease: patterns of linkage disequilibrium and disease/marker association. *Genomics* 89, 655–665. doi: 10.1016/j.ygeno.2007.02.002
- Zhang, Q., Sidorenko, J., Couvy-Duchesne, B., Marioni, R. E., Wright, M. J., Goate, A. M., et al. (2020). Risk prediction of late-onset Alzheimer's disease implies an oligogenic architecture. *Nat. Commun.* 11:4799. doi: 10.1038/s41467-020-18534-1
- Zhang, X., Zhu, C., Beecham, G., Vardarajan, B. N., Lancour, D., Farrell, J. J., et al. (2019). A rare missense variant of CASP7 is associated with familial late-onset Alzheimer's disease. *Alzheimers Dement.* 15, 441–452. doi: 10.1016/j.jalz.2018.10.005

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2021 Monk, Rajkovic, Petrus, Rajkovic, Gaasterland and Malinow. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Chapter 1, in full, is a reprint of the material as it appears in *Frontiers in Genetics*. Monk, Bradley; Rajkovic, Andrei; Petrus, Semar; Rajkovic, Aleks; Gaasterland, Terry; Malinow, Roberto, 2021. The dissertation author was the primary investigator and author of this material.

CHAPTER 2

Synaptic stability and plasticity: a role of the actin filament network

Monk, B., Dore, K., Proulx, C., Alphonso, S., Marino, M., Aronson, S., Malinow, R.M.

Manuscript in preparation for publication

ABSTRACT

It is generally accepted that neural networks integrate new information by modifying *synaptic weights*, a term that refers to the signaling efficacy between neurons at a given synapse. To retain new information as long-term memories, synaptic weights must then adopt a state of relative stability. How brief signals induce persisting changes to synaptic weights, and how these weights are maintained for periods far outlasting the lifetime of synaptic molecules is unknown. Here we examine a simple proposal: synaptic weights are maintained despite molecular turnover because they are proportional to actin filament content in synaptic regions. In the simplest case, one can consider a single dimensional filament; its length can be maintained by ‘treadmilling’: individual actin monomers are added on one end while removed on the other end. To explore this idea in more biologically realistic conditions we developed a computational model, and performed various empirical experiments probing actin dynamics in neural dendrites. In simulating plasticity, the model shows that (a) a filamentous actin network can remain stable indefinitely, despite molecular turnover; and (b) transient increase (or decrease) in available actin monomers can rapidly lengthen (or shorten) filaments, which retain their new length after monomer concentration returns to baseline levels. Empirical experiments support essential elements of this model: (a) filament networks in dendritic spines are stable far longer than their individual actin subunits; (b) increasing synaptic filamentous actin increases synaptic weights; and (c) transient signals that release monomeric actin from reserve pools cause spine growth and long-term potentiation.

INTRODUCTION

Memories are thought to be encoded by *synaptic weights*, a term that refers to the strength or efficacy by which an upstream neuron evokes downstream responses via a particular synapse. In many brain regions synaptic weights are largely determined by postsynaptic AMPA-type glutamate receptor (AMPA) levels, which mediate fast-excitatory transmission in the central nervous system (Bassani et al., 2013; Brecht & Nicoll, 2003; Lüscher et al., 1999; Malinow & Malenka, 2002; Song & Huganir, 2002). Given these weights encode a lifetime of memories, synaptic receptor counts must remain stable for many years; yet all synaptic proteins undergo constant turnover (Ehlers et al., 2007; Shi et al., 2001). How synaptic weights remain stable despite continuous molecular turnover is unknown (Shouval, 2005; Smolen et al., 2019). This chapter addresses two fundamental questions on the neurobiology of memory: (a) how synaptic weights are maintained for periods far outlasting the lifetime of synaptic molecules, and (b) how temporary signals induce persisting changes to synaptic weights.

The question of how memories persist orders of magnitude longer than their molecular substrates has been of interest since the 1980s (Crick, 1984). Various attempts to explain this phenomena have centered around theoretical *molecular switches*: molecules that, when activated, remain activated indefinitely (see Figure 2.1A) (J. E. Lisman & Goldring, 1988; Sacktor, 2012; Si & Kandel, 2016). However, the role of molecular switches in memory maintenance remains poorly understood (Jones, 2013; Otmakhov et al., 1997; Volk et al., 2013).

Shouval (2005) proposed an interesting alternative to molecular switches, where LTM maintenance is achieved through receptor clustering. In this computational model, surface receptors form *metastable* clusters governed by local interactions (see Figure 2.1B) - *metastable* meaning the cluster can remain a stable size while all its individual subunits turn over. This model is elegant in its simplicity, addressing both questions (a) and (b) raised above. The synaptic membrane is represented as a uniform

2D grid, where any given lattice position is either vacant or occupied by a receptor. This can be represented as a matrix of zeros (vacant) and ones (occupied):

```

0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0
0 0 1 1 1 1 1 0 0
0 0 1 1 1 0 1 0 0
0 0 1 1 1 1 1 0 0
0 0 0 1 1 1 1 0 0
0 0 1 1 1 1 0 0 0
0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0

```

Receptors occupy and vacate lattice positions in a probabilistic manner. Any given position has an *on-rate* probability that depends on the number of neighboring receptors (immediately above, below, left and right, with 0–4 possible neighbors). The *off-rate* probability does not consider neighbors, and is uniform across the lattice. At the crux of this model is (1) a neighbor dependent on-rate that provides a means to prevent cluster growth, while (2) a uniform off-rate ensures total receptor turnover. The on-rate formulation to prevent cluster growth is particularly innovative, circumventing the need to actively change rate parameters to achieve metastability. Clusters are sustained because $k_{on} \gg k_{off}$ within the cluster (2, 3, or 4 neighbors), and because $k_{on} \approx 0$ outside the cluster (1 or 0 neighbors).

A primary issue with the Shouval cluster model is that it lacks biophysical and biological support. It is generally thought that receptors are not exocytosed at the postsynaptic membrane, as would be necessary to fill gaps in the center of the cluster; they are instead inserted into dendritic membrane regions outside of spines and diffuse laterally along the surface into spines and postsynaptic areas (Borgdorff & Choquet, 2002; Makino & Malinow, 2009; Yudowski et al., 2007). Also there is no evidence that AMPARs directly interact with each other to form clusters (Bassani et al., 2013). Another issue is that this model requires that *on-rate* probability for two, three, and four neighbors is relatively high ($p \approx 0.1$), while for one neighbor it is extremely small ($p \approx 0.1 \times 10^{-11}$), which lacks a biophysical basis. It would be more biologically plausible to have a neighbor dependent *off-rate* (i.e. breaking 4 bonds takes longer than 1 bond) and a fixed on-rate that scales with molecular concentration. However under

Shouval's formulation, such clusters are unstable. For these reasons the Shouval cluster model, while an interesting concept that can provide a mathematical insight into the questions (a) and (b) introduced in the first paragraph, the model lacks a biological analog.

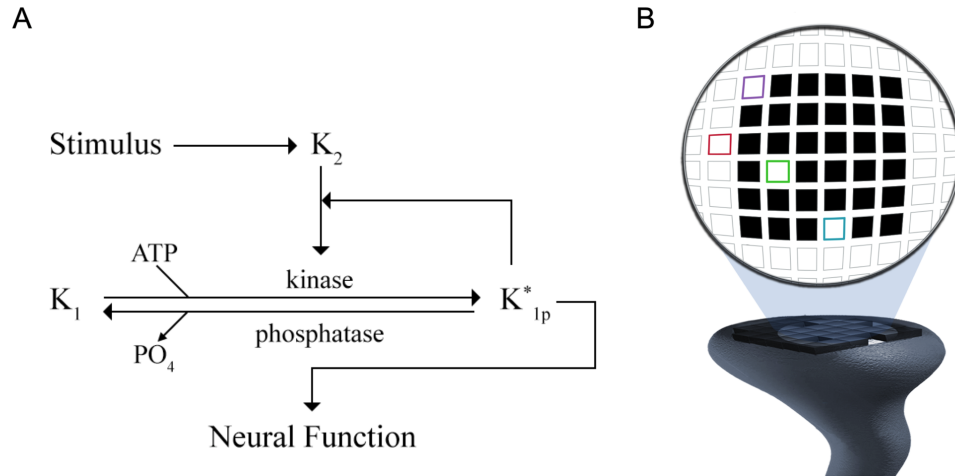


Figure 2.1: LTM models. (A) Depiction of a Lisman molecular switch model. This model shows the primary reactions in a proposed bistable switch. Two kinase proteins, denoted here as K1 and K2, compose the switch. The K1 protein exists in either an inactive or active (K1*) state. Theoretically this reaction can be initiated by neural stimulation, at which point it can be self-sustained. (B) Depiction of the Shouval cluster model of synaptic receptors, projected onto the postsynaptic region of a dendritic spine. In this model surface receptors cluster in synaptic regions with a uniform off-rate and a neighbor-dependent on-rate. Lattice locations highlighted in red, purple, blue, and green have one, two, three and four neighbors, respectively. In this model exocytosis (cluster addition) is neighbor-dependent, while endocytosis (cluster removal) happens at a fixed-rate uniformly across the surface.

Where do we go from here? Our approach towards addressing the two questions above has been to examine the behavior of AMPAR in-and-around synapses, as described in experimental studies (see below), and then propose a metastable mechanism to control their synaptic levels. Studies indicate that AMPARs are inserted into the dendritic membrane outside of spines (Borgdorff & Choquet, 2002; Collingridge et al., 2004; Makino & Malinow, 2009; Yudowski et al., 2007). Once on the surface, AMPARs stochastically diffuse along the membrane where they enter and exit dendritic spines and synaptic areas (Hoze et al., 2012; Nair et al., 2013; Renner et al., 2009). An emerging theory is that

synaptic membrane viscosity or diffusional trapping could dictate synaptic receptor levels (Czöndör et al., 2012; Ehlers et al., 2007; Holcman & Triller, 2006). The idea is simple: as synapses reduce their diffusion rate they increase receptor levels. A similar, and likewise parsimonious idea is that as spines and synapses increase in size they contain more receptors. Indeed, LTP signals have been shown to induce dendritic spine growth (Fischer et al., 1998; Kopec et al., 2007; Lang et al., 2004).

Both ideas are eloquent solutions to a seemingly complex problem. Though, the key questions mentioned above, (a) and (b), remain. Here we present a model inspired by Shouval's approach to these questions (i.e. identify a metastable system or structure in dendritic spines), but is also supported by physiological evidence. Central to this model is the protein actin, a structural molecule identified to have metastable properties in its filamentous form. Dendritic spine size and synaptic area are generally proportional to a spine's total actin filament content (Honkura et al., 2008; Korobova & Svitkina, 2010). Actin filaments also provide scaffolding for scaffold-associated proteins (SAP) known to interact with synaptic AMPARs (Kessels et al., 2009; Shen et al., 2000) likely reducing their synaptic diffusion rate (Czöndör et al., 2012; Ehlers et al., 2007; Holcman & Triller, 2006; Simon et al., 2013).

Given that actin filaments have metastable properties and spine actin content can influence synaptic AMPAR levels, actin could be a primary regulator of synaptic weights. Our findings support this conclusion - that the metastable properties of actin filaments allow synaptic weights to be modified by transient signals and attain long-term stability despite total molecular turnover. These findings are supported by both computational simulations and biological experiments that demonstrate how actin dynamics explains (*a*) how are synaptic sizes and/or diffusion rates maintained for periods far outlasting the lifetime of synaptic molecules, and (*b*) how do temporary signals induce persisting changes to synaptic sizes and/or diffusion rates.

RESULTS

PART I: COMPUTATIONAL MODELS

Dendritic spine area & diffusion rate can determine synaptic receptor levels: AMPARs diffuse stochastically along the membrane where they enter and exit dendritic spines and synaptic areas. Using a computational model, we tested whether synaptic membrane viscosity could stably mediate synaptic receptor levels. Receptor surface diffusion was first simulated using a simplified scaled model of a $3 \mu\text{m} \times 6 \mu\text{m}$ dendritic segment with two $0.8 \mu\text{m}$ diameter synaptic regions (Figure 2.2A). Parameters were set to empirically measured receptor diffusion rates (Ehlers et al., 2007) across extrasynaptic dendrite surfaces (e.g. $0.15 \mu\text{m}^2/\text{s}$) and postsynaptic regions (e.g. $0.01 \mu\text{m}^2/\text{s}$).

Simulations demonstrated that although receptors are diffusing randomly in-and-out of various dendritic compartments, synaptic receptor density is remarkably stable for hundreds of minutes during baseline conditions (Figure 2.2B-C). We also find that changing the diffusing rate of a particular synapse is sufficient to stably increase or decrease surface receptor counts in the respective synaptic region. In general a region with half the diffusion rate (e.g. $0.01 \mu\text{m}^2\text{s}^{-1}$ vs. $0.02 \mu\text{m}^2\text{s}^{-1}$) of another similarly-sized region will have twice the receptor density; a region with one-third the diffusion rate of another will have a three-fold higher receptor density, etc. (Figure 2.2D).

Simulations also demonstrated that changes in synaptic area can parallel the effects of diffusion rate with regard to synaptic receptor counts (i.e. *total receptors* as opposed to *receptor density*). For example total receptor count will increase by precisely the same amount in a given synapse, whether it doubles its area or halves its diffusion rate. (Figure 2.2E-F). A key takeaway from these computational simulations is that synaptic area and diffusion rate can mediate synaptic receptor levels, which remain stable for periods that outlast the lifetime (i.e. dwell time) of individual molecules. Changes to either synaptic area or synaptic diffusion rate will have a similar effect on total receptor counts, while only diffusion rate, not synaptic area, can modulate synaptic receptor density.

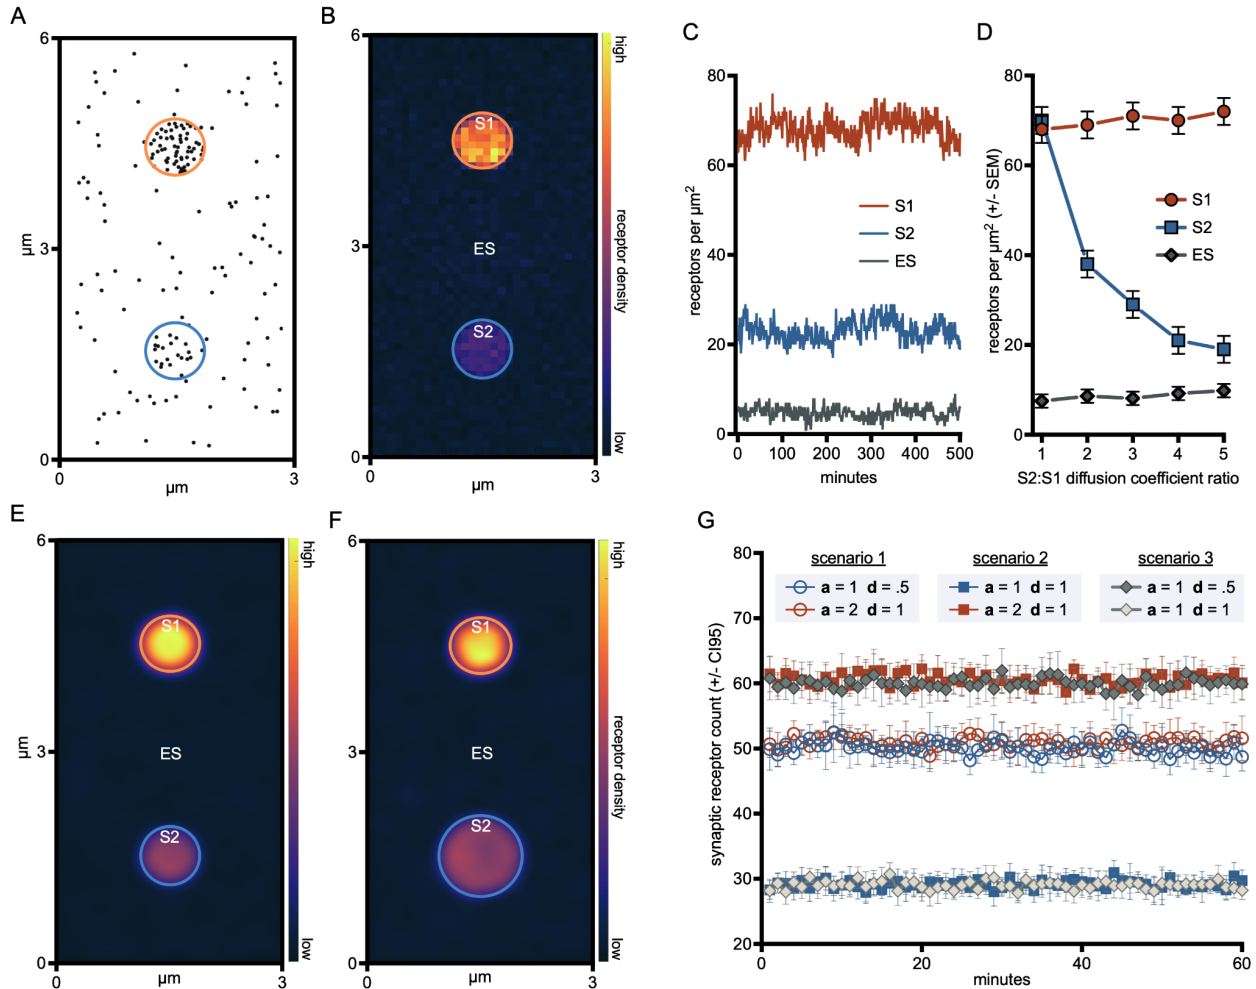


Figure 2.2: Receptor density is a function of local diffusion rate. Simulating receptor surface diffusion using a simplified scaled model of a $3 \mu\text{m} \times 6 \mu\text{m}$ dendritic segment with two $0.8 \mu\text{m}$ diameter synaptic regions. Global diffusion rate was set to $0.1 \mu\text{m}^2/\text{s}$ while reduced diffusion rates were tested in synaptic regions. (A) Left-panel shows a snapshot of stochastic receptor locations after reaching a diffusional steady-state; diffusion in the upper synaptic area (*S1*) was set to $0.01 \mu\text{m}^2/\text{s}$, while the lower synaptic area (*S2*) was $0.02 \mu\text{m}^2/\text{s}$. (B) Heatmap of mean receptor density over a 60-minute window using the same conditions as A. (C) Receptor densities are remarkably stable during steady state turnover. At diffusion rates in A receptor densities were stable for hundreds of minutes. (D) Surface diffusion rate is proportional to receptor density, such that a region with half the diffusion rate of a nearby region sharing a common pool of receptors will have twice the receptor density; a region with one-third the diffusion will have a three-fold higher receptor density, etc. (E) Heatmap of mean receptor density when *S1* has $\frac{1}{2}$ diffusion rate of a similarly sized synapse, *S2*. The result is that *S1* has twice the receptor density and twice the total number of receptors. (F) Heatmap of mean receptor density when *S1* has $\frac{1}{2}$ diffusion rate and $\frac{1}{2}$ the area of *S2*. The result is that *S2* still has half the receptor density, but the same total number of receptors as *S1*. (G) Three scenarios were independently simulated, where synaptic areas *S1* and *S2* were assigned an area ‘*a*’ and diffusion rate ‘*d*’. These simulations further demonstrate that changes in synaptic area and diffusion rate have proportional effects on total receptor counts. However only changes to diffusion rate, not area, impact synaptic receptor density.

The case for actin as a primary regulator of synaptic weights: Since surface area and diffusion rate can dictate synaptic receptor levels, we sought to identify synaptic proteins that could influence synaptic membrane viscosity or dendritic spine size. We examined the properties of actin, as the amount of this molecule in a spine correlates with its size (Cingolani & Goda, 2008; Hotulainen & Hoogenraad, 2010) and actin can bind to proteins that could affect synaptic membrane viscosity (Kusumi et al., 2011; Renner et al., 2009)(Bassani et al., 2013; Derkach et al., 2007; Lin & Webb, 2009; Opazo et al., 2012; Sainlos et al., 2010). Actin continuously cycles between a diffusible monomeric state (G-actin), and a filamentous state (F-actin) whereby many actin molecules assemble into linear polymers (see Figure 2.3A). As a structural protein highly expressed in dendritic spines, actin may influence synaptic efficacies in several ways. Studies indicate that rapid actin filament polymerization precedes dendritic spine growth during LTP (Lin & Webb, 2009; J. Lisman, 2003); and blocking actin polymerization disrupts LTP (Chen et al., 2007; Fukazawa et al., 2003b; Krucker et al., 2000; Okamoto et al., 2009; Ramachandran & Frey, 2009) . Controlling synaptic size is one way actin may influence synaptic weights; another is through its influence on surface receptor diffusion rate. Actin filaments increase membrane viscosity and molecular drag coefficients (Kusumi et al., 2011; Renner et al., 2009). Furthermore, actin filaments support and localize scaffold-associated proteins (SAP). SAP are multivalent proteins known to form clusters, anchor to actin scaffolding, and interact with a variety of molecular species including surface receptors (Bassani et al., 2013; Derkach et al., 2007; Lin & Webb, 2009; Opazo et al., 2012; Sainlos et al., 2010). SAP clusters ultimately anchor to actin scaffolding, so cellular regions with more filaments can support greater SAP densities. When these densities are localized near synaptic membranes, SAP can bind surface receptors to reduce their diffusion rate. Thus, actin may govern synaptic weights by regulating synaptic size and receptor diffusion rate.

There is another enticing reason to explore actin's role in synaptic plasticity: actin filaments seem to have metastable clustering properties. As a 1-dimensional (1-D) polymer, actin filaments do not suffer from the same '*edge-effects*' that limit the stability of 2-D clusters. The stability of 2-D clusters like those described above (Shouval, 2005) are dependent on cluster/perimeter size; small clusters tend to dissipate

completely, while large clusters tend to display runaway growth. In contrast, a 1-D cluster such as a filament only has two ends, and are not subject to perimeter effects that scale with size. With regard to actin filaments, monomers can only be gained or lost at either end of the filament, regardless of its internal length. This suggests that actin filaments can be stable through a range of different lengths (see Figure 2.3A,B).

Importantly, actin filaments can maintain their length while undergoing complete monomer turnover. Actin filaments are asymmetrical, with a so-called ‘barbed \oplus end’ and a ‘pointed \ominus end’. While both ends are capable of polymerization and depolymerization, the critical concentration (CC) of free actin monomers required for polymerization is lower at the \oplus end than \ominus end. When free G-actin concentration falls between the \oplus/\ominus end CCs (as will during steady-state, details below), monomers are added to the \oplus end, and lost at the \ominus end, at the same rate. As a result, actin fibers maintain their length while undergoing complete subunit turnover. Taken together, actin has several properties that may explain how synaptic weights can remain stable despite constant molecular turnover, and how synaptic weights can stabilize at new levels when evoked by transient signals.

To explore this idea in detail, we developed a computational model of actin polymerization dynamics. Actin polymerization has been extensively studied, quantified, and modeled (for an excellent review see: (Bindschadler, 2010)). Like other models, ours integrates well established rate parameters observed across independent empirical experiments (Bindschadler et al., 2004; Halavatyi et al., 2009; Yarmola et al., 2008). Actin models have been categorized into several broad classes (Bindschadler, 2010), and following this schema our model would be considered a time-dependent, Monte-Carlo Markov-Chain (MCMC), spatially-resolved model (Figure 2.3C). To our knowledge it is the only such model integrating these three aspects, which is mainly attributed to the spatially-resolved feature of our model.

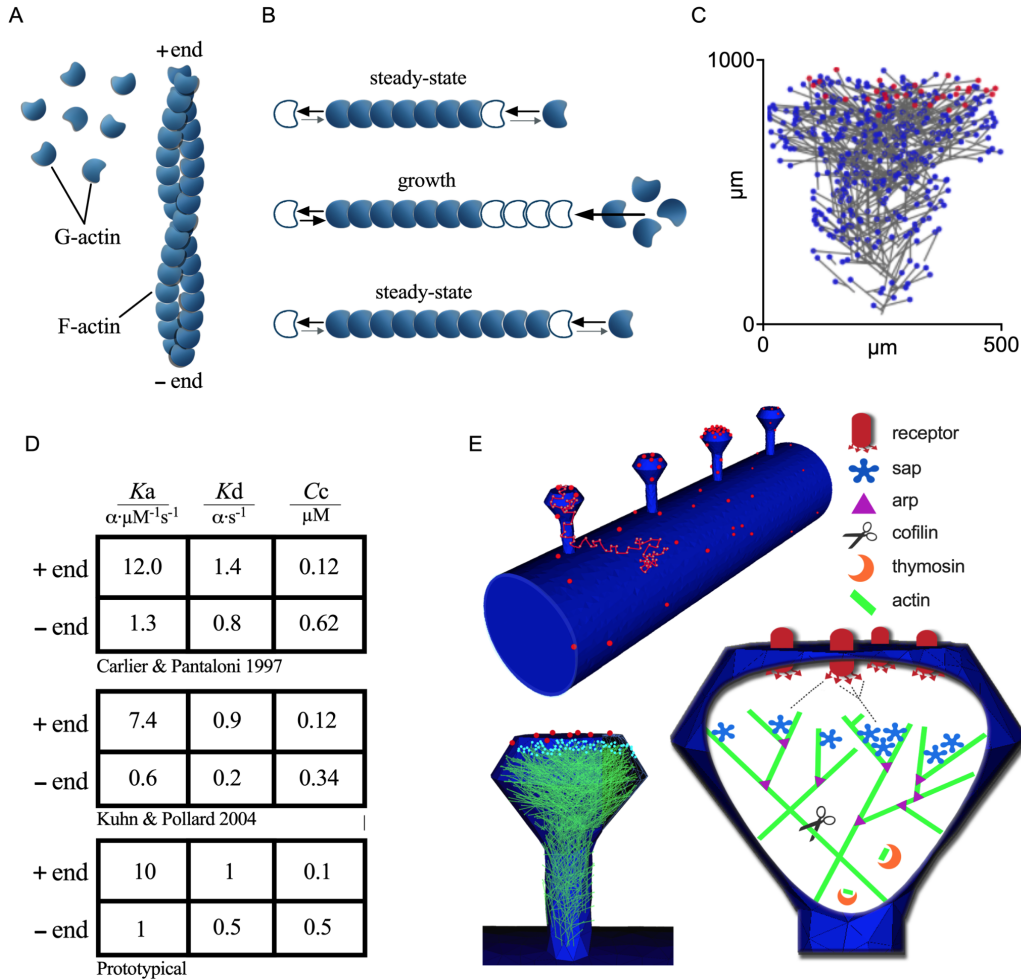


Figure 2.3: Metastable actin synaptic efficacy model. (A) Actin protein can exist as monomeric G-actin or as filamentous F-actin. (B) Actin fibers contain a barbed \oplus end and a pointed \ominus end, each capable of adding and removing actin monomers. The critical concentration required for polymerization is lower at the \oplus end than \ominus end; at steady-state the number of actin subunits added to the \oplus end is proportional to the subunits lost at the \ominus end. If G-actin levels rise, filaments will increase in length until G-actin concentration drops back to steady-state levels. At steady-state actin levels, monomers are lost and added to filaments at equal rates; this exchange rate is independent of filament length. (C) Actin polymerization and branching were simulated in dendritic spines using time-dependent MCMC spatially-resolved modeling. This procedure can be used to generate the spatial coordinates of filament networks in euclidean 3-space as they evolve over time. Shown here is a side-profile view of a dendritic spine filament network that was procedurally generated using our model. (D) Empirical values used to model actin polymerization are shown in the table as prototypical +end and -end polymer on-rates (K_a) and off-rates (K_d). Also shown is the critical concentration (C_c) of G-actin required for net monomer addition at a given end. (E) Actin polymer dynamics play a central role in our model of synaptic plasticity, shown above. Primary components of this model included [1] a dendritic segment with several prototypical spines, [2] diffusible surface receptors, [3] a dynamic actin filament network (including actin, arp, thymosin, cofilin molecules; see text), and [4] multivalent SAP that can interact with filaments, surface receptors, and other SAP. Components of this model can be graphically rendered across all timesteps of the simulation, and viewed as a live animation.

MCMC methods and empirically established parameters were used to model time-evolved actin filament networks inside dendritic spine volumes (Figure 2.3C,D), a centerpiece of our model of synaptic plasticity. In this model, actin serves as the primary coordinator of synaptic weights through its influence of spine size and receptor diffusion rate. In short, the primary components of our model of synaptic plasticity include: (1) receptors that diffuse along the surface of a short dendritic segment with several prototypical spines, (2) a dynamic actin filament network within spine regions, and (3) multivalent SAP that interacts with actin filaments, surface receptors, and other SAP (Figure 2.3E). Since actin dynamics are central to this model, we present those findings in detail next, and subsequently show observations regarding other components of our model.

Simulated actin filament networks are metastable: To determine if synaptic weights can be regulated by interactions between surface receptors, SAP, and actin filaments, we simulated the real-time spatial molecular dynamics of actin filament networks. In addition to actin, the proteins arp2/3 (arp) for filament branching, thymosin- β 4 (thymosin) for G-actin sequestering, and cofilin for filament severing, were included in the model (see METHODS). We find that simulated actin filament networks display long-term stability. Actin subunits rapidly transition between (globular, monomeric) G-actin and (filamentous, polymeric) F-actin, while maintaining stable molecular concentrations (Figure 2.4A-B). Likewise under steady-state conditions, the average length across all filaments (Figure 2.4C), as well as the total number of filament branches found throughout the entire spine filament network (Figure 2.4D) remain stable for many simulated hours.

We next examined whether all F-actin undergoes turnover in the context of a branched filament network. To address this, we evolved a series of actin networks until they reached a steady-state with regard to overall morphology. After these networks settled into a morphological steady-state, the simulation was briefly paused and each filament in the network was tagged with a unique identifier. The simulation then continued, while tracking the lifespans of each tagged filament. In these simulations, most filaments survived less than 1-hour, and no single filament was sustained longer than 5 hours (Figure

2.4E,F). Thus it appears that metastability in the context of a multi-filament network is represented at the network-level: total F-actin content, average filament length, and number of filament branches within the network remain stable while individual monomers turn-over within filaments, and individual filaments turn-over within the network.

We next tested whether average filament length within an actin network stably increases following transient growth signals. As mentioned above, thymosin sequesters G-actin through reversible binding, which creates a large reserve pool of G-actin (Xue & Robinson, 2013), that could be released from thymosin by a transient signal. During steady-state conditions, the concentration of *free* G-actin (i.e. not bound by either thymosin or filaments) remains stable (Figure 2.4G). At the 2-hour mark, we transiently reduce the affinity between thymosin and G-actin, which liberates many sequestered actin molecules into the pool of *free* G-actin. These monomers are quickly absorbed by the filament network, significantly and stably increasing filament lengths throughout the network (Figure 2.4H,I). Interestingly, this transient influx of monomers also increases, stably, the total number of network branches (Figure 2.4J).

This outcome demonstrates that branch formation rate is a function of the total length of all filaments within a network. Arp protein can bind anywhere along a filament to create a new branch; therefore, as the total length of a filament network increases, the probability-rate of arp binding also increases (and as soon as arp binds a filament, it nucleates a new branch site). Lastly, we observe that the increased mean filament length and increased branch count both remained stable, despite all F-actin content undergoing complete turnover (Figure 2.4K,L). The time constant for F-actin turnover was not significantly different pre- and post-growth.

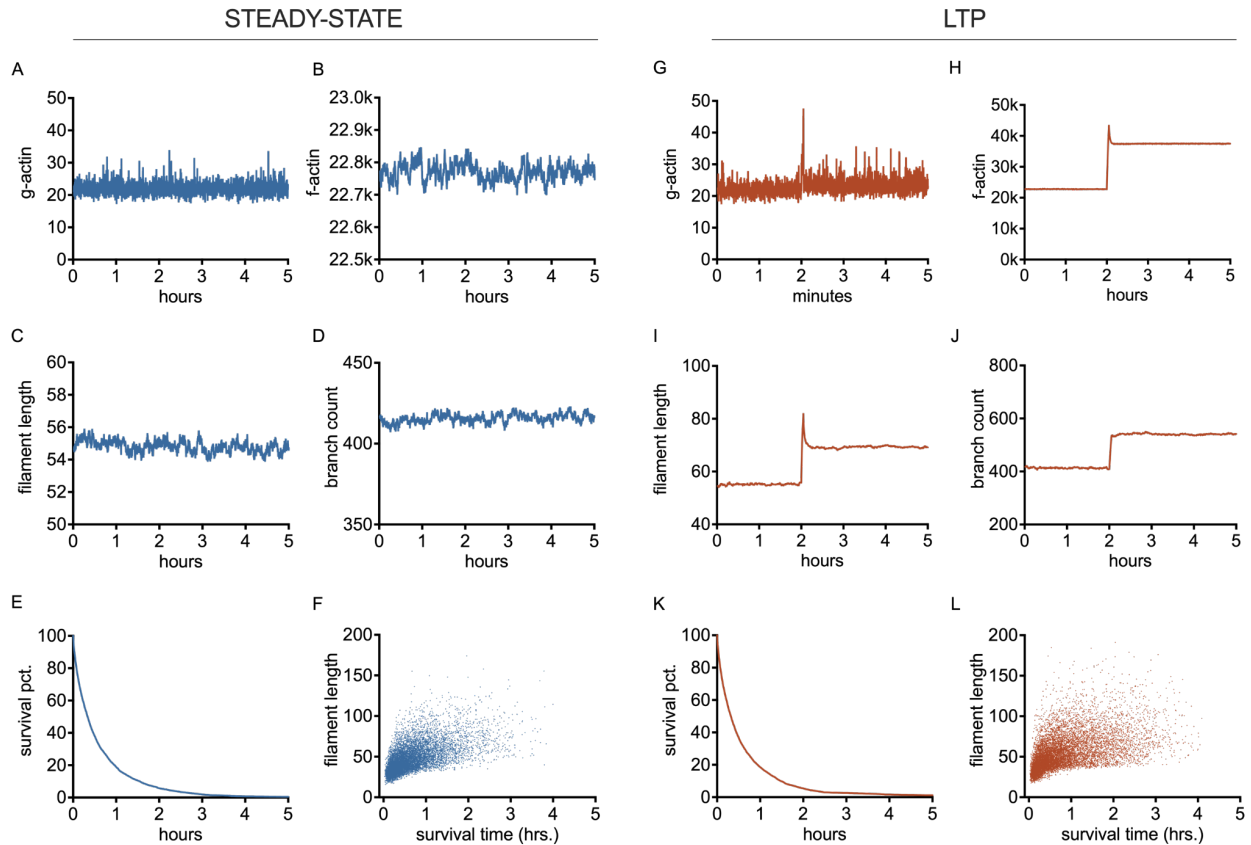


Figure 2.4: Simulated actin dynamics inside dendritic spines. All plots in this figure use data averaged over 10 independent runs, with a simulated real-world time of 5 hours. The x-axis of all plots represent time; and all display exactly the same 5-hour (300-minute) time window. Each metric plotted on the y-axis refers to a quantity measured inside a single dendritic spine. (A) Diffusible unbound G-actin monomers. (B) The total number of F-actin subunits composing all filaments inside a spine. (C) Average number of F-actin subunits composing each filament. (D) Total number of filaments; equal to the total number of branches across the entire spine filament network (E) The spine filament network was allowed to evolve until reaching a steady-state, at which point each filament in the network was tagged with a unique identifier. The filament network was then allowed to continue evolving, and the lifespans of the tagged filaments were tracked. The plot here shows the percent of tagged filaments remaining as the network evolved over 5 hours. Most filaments survive less than 1 hour; no filament from any run lasted the full 5 hours. (F) Survival times as a function of filament length. (G-L) Same as A-F, except a brief LTP event was delivered at the 2-hour (120-min.) mark. This event consisted of temporarily making a reserve pool of G-actin available for polymerization. This reserve pool was made available for 30 sec. after which point G-actin levels were returned to baseline.

Metastability of actin filament networks carries through to synaptic receptor levels: Actin filament network dynamics, SAP clustering, and receptor diffusion were modeled as three layers of interactions (see METHODS). Filament networks were simulated inside dendritic spine regions, with quantitatively

similar outcomes as those described above (see Figure 2.4). In this model, filaments located near postsynaptic surfaces acted as seed points for SAP clusters (Figure 2.5A). Clustering behavior was simulated using a customized neighbor-dependent algorithm (see above and METHODS). The primary difference between our algorithm and Shouval (2005): in our model lattice on-rate probability is a function of actin filament proximity, and lattice off-rate is neighbor dependent. These updates to Shouval’s clustering algorithm better reflect an expected biophysical basis.

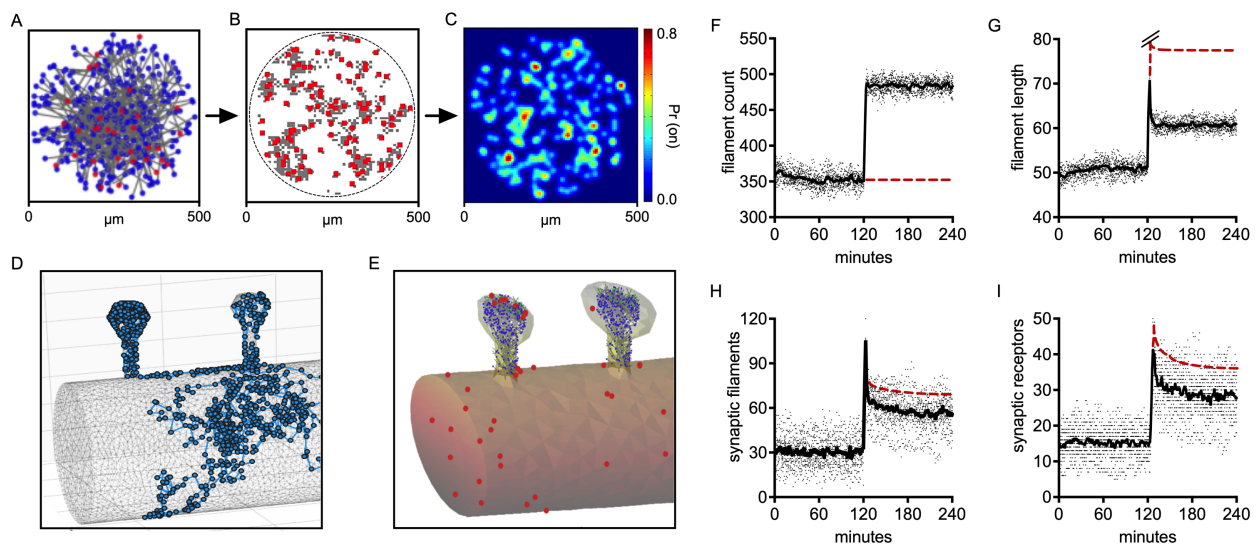


Figure 2.5: Simulated actin dynamics inside dendritic spines and receptor count at synapses. An MCMC spatially resolved model was used to simulate actin filament dynamics inside dendritic spine regions. **(A)** Overhead view of procedurally generated filament network inside prototypical spine. Filament that fell within 20 nm of postsynaptic membrane (PSD) colored red, otherwise colored blue. **(B)** Filaments in PSD acted as seed-points for SAP clusters, simulated using a neighbor-dependent algorithm where on-rate was a function of actin filament proximity, and lattice off-rate was neighbor dependent. Filament tips shown in red; SAP in gray. **(C)** Heatmap displays SAP-receptor interaction probabilities. Synaptic locations highly populated by SAP have high interaction probabilities compared to regions with sparse SAP. Probabilities dynamically change each timestep according to stochastic behavior of proteins. **(D)** Surface receptor diffusion was simulated along a prototypical dendritic segment with several spines. Shown here is a 10-min trace of a single receptor, along the dendrite shaft and both spine regions. **(E)** Integrated model with all components simulated in parallel. **(F-I)** Independent simulations were performed that tracked filament count, mean filament length, synaptic filament count, and synaptic receptors during a 120 min. baseline. At 120 min. a brief 60 sec. LTP signal was delivered. Same metrics were tracked for another 120 min. In each plot, the solid line is mean of all runs; dots are values from each replicate. Red dashed lines are mean outcomes from replicates where filament network was not allowed to exceed 350 total branches. **(F)** Total number of filaments in each actin filament network. **(G)** Mean length across all filaments in each actin network. **(H)** Number of actin filaments at each timepoint located inside the PSD (<20 nm from the postsynaptic membrane). **(I)** Number of surface receptors in synaptic regions.

In our model the precise spatial location of every actin filament (Figure 2.5A), SAP cluster molecule (Figure 2.5B), and surface receptor (Figure 2.5D) is known. This allows us to compute proximity-dependent binding probabilities to model the interactions between surface receptor and SAP (Figure 2.5C). Each interaction made between a SAP molecule and a surface receptor was transient, and consisted of briefly reducing the receptor's diffusion rate. The reduction of diffusion rate during a transient interaction was based on empirically measured values (Ehlers et al., 2007)

Using all components of the model (Figure 2.5E) (filament dynamics, SAP clustering, and surface receptor diffusion) in parallel, we performed ten independent simulations representing 4-hours each of real-world time. In these simulations we tracked filament counts, average filament length, synaptic filament counts, and synaptic receptors throughout the entire simulation. The first 120-minutes consisted of a baseline period, to evaluate steady-state dynamics. To mimic an LTP signal, at the 120-minute mark, we transiently reduced the thymosin-actin binding for 30-seconds; after which point the binding affinity was restored to baseline. Data was again collected for an additional 120-minutes to monitor the effects of the LTP signal.

We find that all metrics are stable during the steady-state period, including the total number of branches composing the filament network (Figure 2.5F), average filament length (Figure 2.5G), number of filaments inside the postsynaptic density (PSD) (Figure 2.5H), and number of surface receptors located in the synaptic area (Figure 2.5I). During LTP we find that all metrics related to the actin filament network abruptly increase, including the number of filaments, average filaments length, and the number of PSD filaments. These changes are immediately followed by a sharp increase in the number of synaptic receptors. To determine whether the LTP-induced increase in total filament branches is necessary to significantly and stably increase synaptic receptors, we performed another 10 runs using the exact same parameters, except total filament branch count was capped at 350 total branches (see dashed red line in Figure 2.5F). Interestingly we find that synaptic receptors levels still increased, and to a greater extent than before branching was limited. This suggests that stable increases in synaptic receptor levels are not

strictly dependent on the number of branches in a filament network; lengthening of existing branches is sufficient to achieve stable increases in receptor counts.

PART II: EXPERIMENTAL TESTS OF KEY MODEL COMPONENTS

Altering actin bioavailability can change otherwise stable dendritic spine morphologies: To test whether dendritic spines in biological neurons generally maintained stable sizes, we expressed GFP-tagged actin in mouse hippocampal slice cultures, and monitored the size of dendritic spines for 1-hour using two-photon laser scanning microscopy. Spine area was computed by manually tracing the boundary around each spine, at each timepoint (Figure 2.6A). The average size across all spines was $0.499 \mu\text{m}^2$ ($n = 20$, $sd = 0.014$). Over the 1-hour monitoring period the average absolute change in size was $0.017 \mu\text{m}^2$ ($sd = 0.016$), with the smallest change being $<0.001 \mu\text{m}^2$ ($\sim 0\%$) while the largest change was measured to be $0.065 \mu\text{m}^2$ (15.7%). Only 3 of the 20 spines changed their relative size ranking over the 1-hour experiment, with the largest rank-change being just 2 positions (from 13th largest to 11th largest). Thus among this sample, relatively small spines stayed small, medium spines remained medium, and large spines persisted being large, which suggests that irrespective of size, dendritic spines can remain stable for at least duration on the order of hours.

One possible explanation for the stability in spine sizes as measured in the experiment above, could be that under experimental conditions, morphological change is diminished due to the slowed or fully-arrested dynamics of the underlying structural molecules. To determine if this was the case, we again expressed GFP-tagged actin in hippocampal primary cultures, and then assessed the turnover rate of actin protein inside dendritic spines through fluorescence recovery after photobleaching (FRAP) experiments. This is a well established technique used to estimate the rate of molecular turning-over within a cellular compartment. Using FRAP, we find that within 60-minutes, tagged actin underwent complete turnover (Figure 2.6B,C).

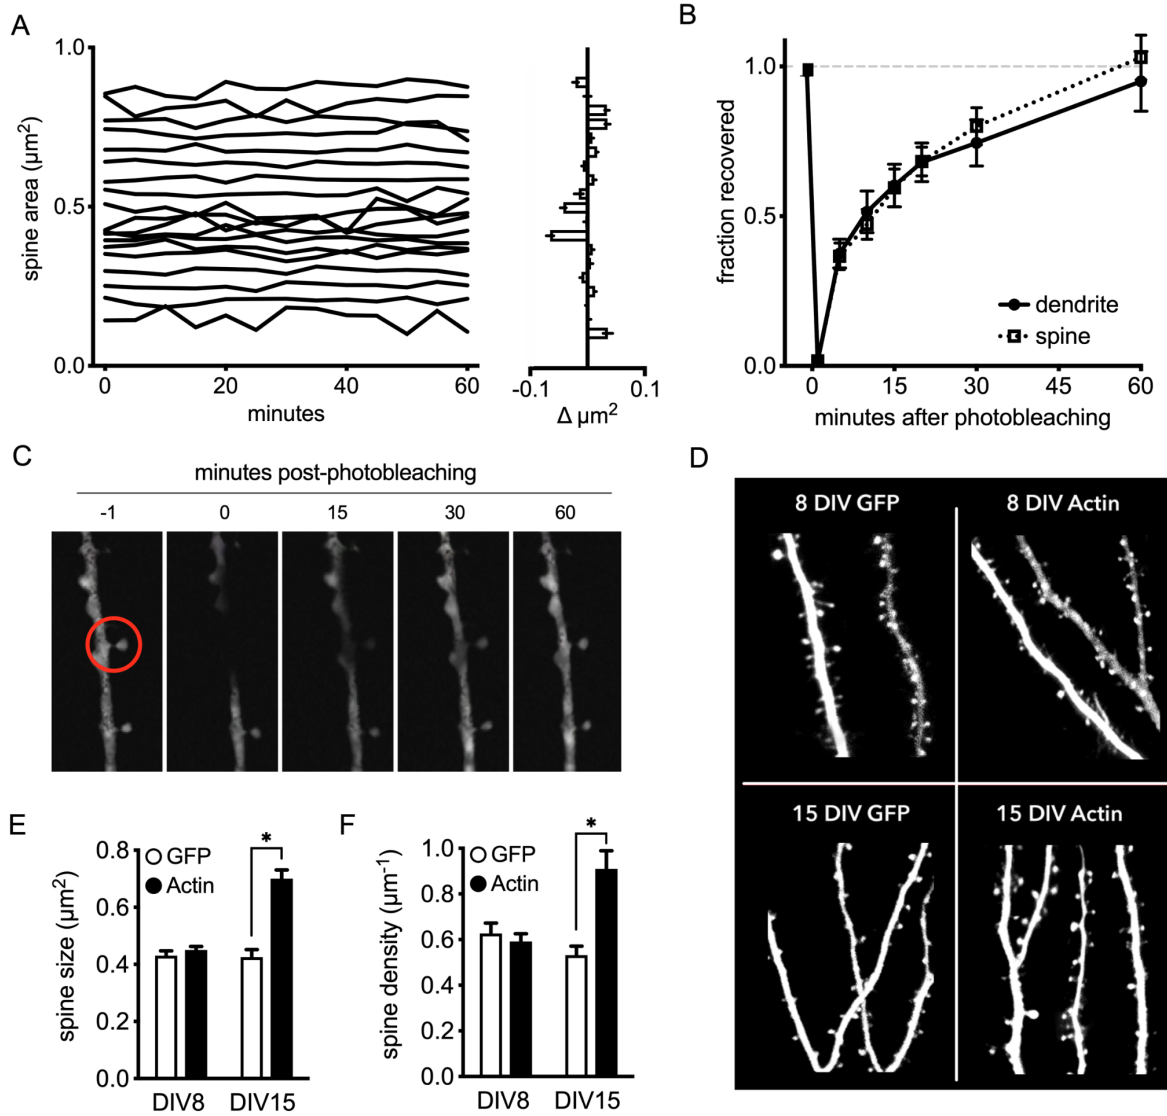


Figure 2.6. Dendritic spines are morphologically stable despite complete actin turnover. (A) To estimate stability of the actin filament network, we monitored the size of dendritic spines expressing GFP-tagged actin with two-photon laser scanning microscopy in hippocampal primary cultures. Irrespective of size, spines remained stable over a 60-minute monitoring period. (B) To measure the turnover of actin we performed FRAP experiments. Images of neuron expressing GFP-tagged actin were collected 1 minute prior to bleaching, immediately after bleaching, then 15 minutes after, 30 minutes after, and 60 minutes after bleaching. (C) At 60 minutes post-photobleaching, actin-GFP fluorescence generally displayed complete recovery, indicating that all actin content in dendritic spines, including filamentous actin, is recycled within approximately one hour. (D) We performed overexpression experiments to determine if the actin overexpression has an effect on dendritic spine size. Images of neuron overexpressing GFP alone or GFP-tagged actin were taken after 8 DIV and 15 DIV. (E) After 15 days DIV, dendritic spines are significantly larger in neurons overexpressing actin than in control neurons. (F) After 15 days DIV, the density of spines along the length of dendrites is significantly greater in neurons overexpressing actin than control neurons. Error bars represent SEM. * $p < .05$, ** $p < .01$

We conclude that fluorescence recovery ($t_{1/2} \sim 10$ min) is rate-limited by filamentous F-actin-GFP in spines, as monomeric G-actin-GFP is expected to recycle in-and-out of spines within seconds (Biess et al., 2007; Holcman & Schuss, 2011; Svoboda et al., 1996). These findings rule out that spine size stability is due to arrested structural protein dynamics; and in fact demonstrate that spines retain their size despite the complete turnover of structural F-actin protein.

To determine whether increased actin availability is sufficient to increase dendritic spine sizes we performed an actin overexpression experiment. On DIV4, hippocampal cultures were incubated with either actin-GFP or GFP alone. Morphological assessments were then performed on DIV8 and and DIV15 (Figure 2.6D). On DIV8 there were no statistical differences between the GFP-actin and GFP control group, with regard to spine size or density. By DIV15 however, we find that hippocampal neurons expressing actin-GFP had significantly larger dendritic spines than neurons only expressing GFP (by 0.28 ± 0.04 microns), $t(161) = 6.74$, $p < .001$ (Figure 2.6E). Likewise, neurons overexpressing actin had a higher spine density than controls, $t(161) = 4.05$, $p < .001$ (Figure 2.6F). These findings suggest the bioavailability of actin can drive significant morphological effects in dendritic spines after DIV8 (there may be some other limiting molecules prior to DIV8). Specifically, we demonstrate that global increases in actin protein expression in neurons results in dendrites with significantly more spines, and significantly larger spines.

Actin overexpression increases synaptic weights: Results from our simulation experiments, along with evidence from previous research (Fukazawa et al., 2003a; Krucker et al., 2000), suggest that increased actin bioavailability in dendrites may be sufficient to increase synaptic weights. To experimentally determine whether actin bioavailability can impact synaptic weights, we performed a series of overexpression experiments and measured spontaneous miniature excitatory postsynaptic currents (mEPSC), which are indicative of functional synapses and their synaptic weight (Ehrlich et al., 2007; Otmakhov et al., 1993; Raastad et al., 1992; Simoni et al., 2003; Wyllie et al., 1994). It is possible that actin overexpression may increase spine size without producing a complementary increase in synaptic receptors. On the other hand, if we find increased mEPSC amplitudes following acute actin

overexpression, it would suggest that increasing actin bioavailability is sufficient to upregulate excitatory receptors in hippocampal synapses.

Here we used Sindbis virus (Makino & Malinow, 2009) to acutely overexpress actin in organotypic hippocampal slice cultures. Within 24 hours of virus exposure, electrophysiological recordings were performed in voltage clamp to measure mEPSC amplitude and frequency in control neurons and in neurons overexpressing actin. Indeed we find a significant increase in the mEPSC amplitude in neurons expressing actin (28.6 +/- 1.4 SEM) compared to control neurons (19.1 +/- 1.2 SEM), $t(43) = 6.74$, $p < .001$ (Figure 2.7A-D). This signifies there are substantially more AMPA receptors in synapses of actin-overexpressing neurons than in control synapses. Interestingly, despite our finding that actin overexpression increases spine density, we find no statistical difference in mEPSC frequency ($p = .09$) (Figure 2.7E).

This could mean one of several things. First, it could be that our dataset is underpowered, and additional data would reveal small group differences in frequency exist. It could also mean that increased actin bioavailability first impacts the sizes of existing spines before evoking *de novo* synthesis of new spines (i.e. high actin levels must be sustained for longer than 24 hours). If however actin expression did acutely increase spine density, this result could mean that some spines are completely starved of AMPA receptors (i.e. 'silent synapses') (Liao et al., 1995; Poncer & Malinow, 2001). While those scenarios will be interesting to explore in future research, the primary result here is that actin overexpression substantially increased the amplitude of synaptic responses.

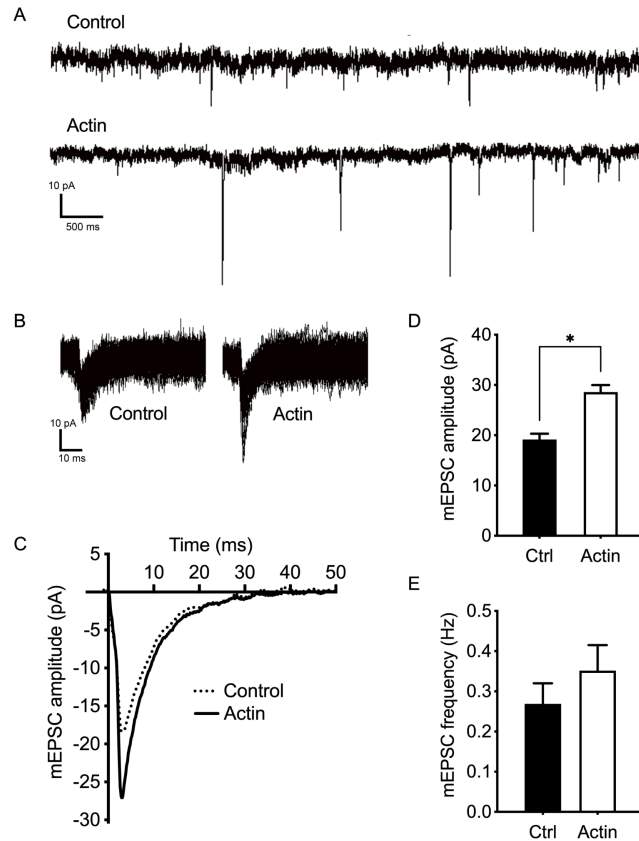


Figure 2.7. Actin overexpression increases excitatory postsynaptic currents. (A) Example of postsynaptic current recordings from control neurons vs. neurons overexpressing actin protein. (B) Stacked overlay of dozens of mEPSCs from neurons overexpressing actin and control neurons. (C) Average shape of spontaneous mEPSCs from neurons overexpressing actin and control neurons. (D) Average mEPSC peak amplitude in neurons overexpressing actin and control neurons. (E) Average mEPSC frequency in neurons overexpressing actin and control neurons. Error bars represent SEM. * $p < 0.001$

Actin dissociates from actin-sequestering protein, thymosin, during LTP: The experiments so far illustrate that upregulating actin availability can drive increases in spine sizes and synaptic weights, in preparations where actin was increased chronically, throughout the neuron, over the course of hours or days. In this final series of experiments we investigate whether acute, local increases in actin bioavailability yields similar effects to those observed during chronic global actin increases. We also assess a specific mechanism that neurons may use to quickly and transiently increase actin bioavailability: the liberation of G-actin from a reserve pool of actin monomers bound by thymosin protein (Figure 2.8A). We reasoned that such an event could drive spine growth during relatively brief LTP signals. If indeed a

significant reduction in thymosin-actin interactions were observed during potentiation signals, it would provide striking insight regarding the structural plasticity component of LTP.

Since relatively little is known about thymosin-actin sequestration, particularly in neurons, we first tested whether thymosin expression had any effect on dendritic spine sizes, on its own, and in conjunction with actin overexpression. Interestingly, we find that thymosin expression on its own does not impact spine size, but does significantly attenuate spine size in actin-overexpressing neurons (Figure 2.8B). That is, aside from the control vs. thymosin comparison, all other pairwise comparisons were significant. Spines in neurons overexpressing actin were significantly larger than control neurons $t(89) = 9.96, p < .001$, and as predicted thymosin attenuated this effect $t(130) = 5.52, p < .001$. However, this was only a partial attenuation, as thymosin+actin overexpressing samples still had larger spines than controls on average, $t(131) = 6.56, p < .001$.

We next used an occlusion-type experiment to test if the increase in transmission produced by over-expressed actin maximized a signaling mechanism used in LTP. If such were the case, one would expect reduced or absent LTP in neurons overexpressing actin. Here we used a chemical-LTP (cLTP) protocol to acutely and globally induce LTP across all neurons in the preparation (see METHODS). We find that before LTP, spines from actin-overexpressing neurons are larger than control counterparts (consistent with results above). The cLTP protocol induced spine growth in both the control group (increase $\mu = 0.12 \mu\text{m}^2, \sigma = 0.04$), $t(19) = 12.9, p < .001$, and to a slightly lesser extent in the actin group (increase $\mu = 0.10 \mu\text{m}^2, \sigma = 0.05$), $t(19) = 56.5, p < .001$ (Figure 2.8C). With regard to fold-change, spines from control neurons displayed significantly greater change than spines from actin expressing neurons, $t(38) = 5.23, p < .001$ (Figure 2.8D).

We also examined whether pre-cLTP spine size had any effect on spine growth. To test this we split spines into equal groups based on original spine size (large and small spines). We used a linear regression model where fold-change was the dependent variable, and the independent variables were spine size, and condition (actin vs. control). We find the overall prediction model is significant, $F(4,36) = 20.6, p < .001, R^2 = .632$, with condition being a significant predictor ($\beta = 0.20; p = .022$), and the

condition-by-size interaction being a significant predictor ($\beta=0.35$; $p=.004$). This suggests that small spines particularly in the control group (that had the smallest spines overall) display the greatest growth relative to their own starting (pre-cLTP) size. These results are consistent with the view that in large spines, increasing actin maximizes a signaling mechanism used in LTP. In smaller spines, increased actin enhances a signaling mechanism that has been used to a similar amount as that used to produce large spines in control cells. In general, increased levels of actin partially or completely occlude LTP.

In order to determine if LTP signals drive a change in the interactions of actin and thymosin, we developed an experimental procedure involving glutamate uncaging and fluorescence lifetime imaging microscopy (FLIM). Glutamate uncaging allows the precise spatial and temporal triggering of LTP events, localized to single dendritic spines. FLIM can be used to quantify protein-protein interactions, with adequate temporal resolution and sub-micron spatial resolution. By co-expressing thymosin tagged with a red fluorescent protein (mCherry or mApple) along with GFP-tagged actin, we monitor actin thymosin interactions before, during and after LTP at individual spines. Our data indicate that actin reversibly dissociates from thymosin during LTP signaling (Figure 2.8E). The experiment was repeated independently in the lab of collaborator (Y. Hayashi), and the same result was observed (Figure 2.8F): induction of LTP evokes a significant dissociation of actin-thymosin dimers, which peaks within 60 seconds of the signal (relative FRET efficiency $\mu = 1.78 \pm 0.12$ SEM), and returns to baseline within 120 seconds (relative FRET efficiency $\mu = 0.06 \pm 0.09$ SEM). Furthermore, we find this transient dissociation is coupled with a long-term increase in dendritic spine size (limited to only spines at the site of glutamate uncaging). Spines in neurons overexpressing thymosin displayed a greater peak than control neurons (peak thymosin $\mu = 2.55 \pm 0.23$ SEM; peak control $\mu = 1.64 \pm 0.21$ SEM), $t(40) = 161.4$, $p < .001$. Also, spines in neurons overexpressing thymosin had greater sustained growth than spines from control neurons (between 5.5-6.0 hours, thymosin $\mu = 1.90 \pm 0.33$ SEM; peak control $\mu = 1.42 \pm 0.29$ SEM), $t(40) = 114.8$, $p < .001$, that persisted the duration of the experiment (Figure 2.8G). This is consistent with the view that overexpression of thymosin produces a larger thymosin-actin pool at spines, and, through LTP-induced actin-thymosin dissociation, increased actin bioavailability during LTP.

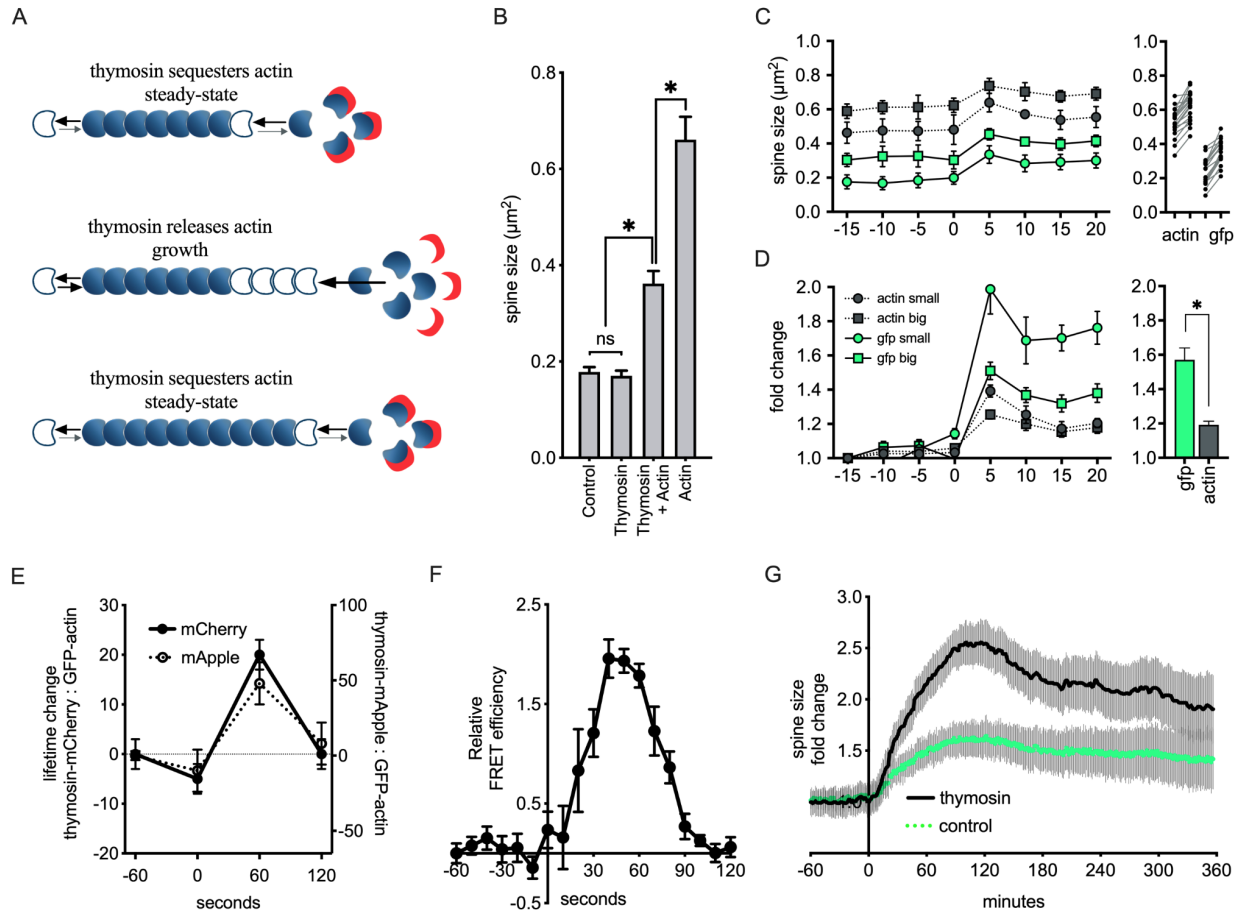


Figure 2.8. Actin dissociates from thymosin-sequestering protein during LTP. (A) During baseline conditions actin (blue) is sequestered by thymosin (red), which creates a reserve pool of monomeric actin. Some actin monomers spontaneously transition from the sequestered reserve pool to the relatively small free polymerizable monomeric pool. An LTP signal temporarily causes thymosin and actin to unbind in high quantities, resulting in filament growth. After the LTP signal dissipates, the system returns to steady-state. (B) Actin overexpression increases dendritic spine area, and thymosin expression partially mitigates this effect. Thymosin expression on its own does not appear to significantly affect spine size. (C) Neurons overexpressing actin had spines that were larger than control neurons, both during baseline, and after cLTP (see plot C for group labels; gray solid and gray dotted lines represent GFP and actin group means, respectively). In general, average spine size of control neurons after cLTP was on-par with actin-overexpressing neurons prior to cLTP. In terms of relative change, control neurons displayed significantly more growth than actin-overexpressing neurons during cLTP. (E) FLIM reveals that actin reversibly dissociates from thymosin during a glutamate uncaging protocol designed to mimic LTP signaling. Two different red fluorophores were used to tag thymosin; both producing similar FRET lifetime change results. Here an increase in lifetime change indicates less thymosin-actin association (F) Similar result as in A, with the experiment repeated independently in another lab. Greater relative FRET efficiency indicates less thymosin-actin association. (G) Average change in dendritic spine size near the site of glutamate uncaging. Spines in neurons overexpressing thymosin displayed significantly greater initial growth than spines from control neurons. This significant size difference persisted for at least the full 6-hour duration of the experiment. Error bars in all figures represent SEM. * $p < .001$

DISCUSSION

Memories are thought to be encoded by synaptic weights, which govern signaling efficacies between neurons. Currently there is no broadly accepted model that explains how synaptic weights are maintained and how temporary signals induce long-term changes to signaling efficacy. In this chapter we propose a model involving actin as the central regulator of synaptic weights, and illustrate through a series of computational and biological experiments how temporary signals induce persisting changes to synaptic actin networks, and how these filament networks are maintained for periods far outlasting the lifetime of synaptic molecules. Our model illustrates how metastable actin filament networks can support both synaptic stability and plasticity, and that synaptic weights could ultimately be a function of the total F-actin content within dendritic spines. These conclusions are supported by both the computational simulations and biological experiments.

The computational modeling experiments in this chapter show that synaptic receptor levels can simply be a function of synaptic diffusion rates. Indeed, despite stochastic turnover of surface receptors, as long as a stable diffusion rate differential is maintained between synaptic areas and membrane regions outside of synapses, a predictable number of surface receptors will accumulate in synapses. However it is unclear how this diffusion rate gradient is achieved, how it is modified, and how it remains stable despite constant protein turnover. To address these questions, we extended our model to include actin, a structural protein highly expressed in synapses. We included actin in our model after identifying its promising metastable properties. It's also known that postsynaptic area and spine size are highly correlated with total actin content within dendritic spines, and actin filaments provide scaffolding for SAP that interact with surface receptors. This last fact links a metastable protein network to surface receptor diffusion rate, thereby providing a potential mechanism for synaptic weight regulation.

The computational modeling experiments incorporating actin show that actin filament networks can maintain average filament lengths despite the complete turnover of their F-actin subunits. Simulation

experiments also show that average filament length within an actin network can stably increase following transient growth signals. Temporarily lowering the affinity between thymosin and G-actin liberates many sequestered actin molecules into the pool of G-actin, which are quickly polymerized into existing filaments. This results in a significant increase in average filament lengths throughout the network. Indeed these newly increased filament lengths remained stable, despite all F-actin content undergoing complete turnover. This stability is maintained even after monomer concentration returns to baseline levels.

Biological experiments also support essential elements of this model. We find that actin filament networks in dendritic spines are morphologically stable orders of magnitude longer than individual G-actin and F-actin subunits. Photobleaching experiments show that all spine actin content is replaced in about one hour, without spines undergoing significant changes in size. This result agrees with our computational models addressing the same phenomena. To assess whether modifying synaptic F-actin content was sufficient to alter synaptic morphology, we globally overexpressed actin in hippocampal neurons. This resulted in significant increases in average dendritic spine size and density.

Next we aimed to determine if upregulating actin bioavailability was sufficient to increase synaptic weights. To test this, we overexpressed actin in organotypic hippocampal slice cultures and performed electrophysiological recordings of mEPSCs. Indeed we found that acute actin overexpression in hippocampal neurons produced a significant increase in mEPSC amplitudes compared to control neurons. This result signifies that acute actin upregulation leads to an accompanying increase in synaptic AMPAR levels. Finally we identify a biological mechanism that can produce acute increases in actin bioavailability at single dendritic spines. We find that under baseline conditions, a peptide called thymosin interacts with G-actin to create a reserve pool of unpolymerized monomeric actin. Furthermore we observe that thymosin releases this reserve pool in response to LTP signals. Indeed this event is followed by a significant and long-lasting increase in dendritic spine size.

Overall, our model provides an example wherein stochastic systems within dendrites (actin, SAP, and surface receptors) can interact and give rise to stability in the aggregate. Through a series of computational and biological experiments, we show that actin filaments have metastable properties and

spine actin content can influence synaptic receptor levels. Together these experiments suggest that actin could be a primary regulator of synaptic weights. Indeed our findings support the conclusion that the metastable properties of actin filaments allow synaptic weights to be modified by transient signals and attain long-term stability despite total molecular turnover.

METHODS

PART I: COMPUTATIONAL MODELS AND SIMULATIONS: All computational models were developed using the Python programming language (Python Software Foundation, version 3.7. available at <http://www.python.org>) and the MATLAB programming and numeric computing platform (R2020a, version 9.8.0, Update 6. The MathWorks Inc.)

Figure 2.2: Receptor diffusion was simulated in MATLAB along a 2-D surface scaled to $3 \mu\text{m} \times 6 \mu\text{m}$, with smaller circular regions (e.g. $0.8 \mu\text{m}$ diameter) representing synaptic areas. To simulate surface diffusion using empirical values requires a decomposition of reported diffusion rate coefficients. A diffusion rate coefficient D (in units: $\mu\text{m}^2/\text{s}$) technically represents to a constant velocity, however it is often used to describe the average velocity of a particle as it moves along a multitude of trajectories over time. If D is known for a given particle under a given set of conditions, it can be implemented in true-to-scale computer simulations of Brownian motion. Concretely, $D = L^2 / (2 m t)$, where L is the step length vector, m is the number of diffusible dimensions, and t is elapsed time.

To accurately simulate diffusion rate to-scale, an important value to calculate is k , the standard deviation of the diffusion rate step size distribution. This value can be determined using the equation: $k = \text{sqrt}(m D)$, where m is the number of diffusible dimensions, and D is the diffusion rate coefficient. The reason this constant is valuable is because it can be used to directly scale random values pulled from a standard normal distribution $N(\mu=0, \sigma=1)$. To demonstrate the utility of this constant, say there is a 2-by-10 array A_t that contains the x and y coordinates for 10 particles at time t ; also there exists a function `randn(2,10)`, that returns a 2-by-10 array of random values from a standard normal distribution. Then the

following equation can be used to update particle locations: $A_{t+1} = k \times randn(2,10) + A_t$. This equation can be used to update particle locations for each and every time step t . to produce Brownian motion at exactly a diffusion rate D , according to Einstein's equations on the theory of Brownian movement (Einstein, 1956).

The generation of all panels of this figure involved simulating particle displacement according to the equations above. Diffusion rates for each surface region are given in the figure caption, or as labels on each figure panel; global diffusion rate was always set to $0.1 \mu\text{m}^2\text{s}^{-1}$. In general the total number of particles used in each simulation was a fixed value (e.g. 400). Boundary conditions were modeled such that particles deflected off the outer edges of the enclosure as if bouncing off a wall. while reduced diffusion rates were tested in synaptic regions. In panel G, instead of giving rate constants in real units of $\mu\text{m}^2\text{s}^{-1}$, diffusion rates are shown as arbitrary units 1 and 0.5 to emphasize the comparison between a given diffusion rate D and $\frac{1}{2}D$.

Figure 2.3: Experimentally determined values (e.g. actin filament K_a and K_d) were used in an MCMC procedure to simulate actin filament polymerization dynamics inside dendritic spines. The rate parameters used in our model can be found in (Bindschadler et al., 2004; Halavatyi et al., 2009; Yarmola et al., 2008). Full simulation code can be found online at: github.com/bradmonk/dissertation. The molecular components used in the filament simulation model are actin, arp, thymosin, cofilin.

Briefly, actin polymerization was simulated by tracking the filament \oplus/\ominus end coordinates in Euclidean 3-space. It's known that each F-actin subunit adds approximately 2.71 nm of length to the filament. Thus the filament origin, along with the number of F-actin subunits in each filament, and the x-, y-, z-plane angle of each filament allows one to solve the coordinates of the filament tip (using 3-D rotational matrices) for each timestep. Knowing the 3-D coordinates of both ends of the filament allows computer software to quickly render the entire filament network at any timepoint.

To simulate thymosin-actin interaction dynamics, we use the Law of Mass Action (LMA). LMA describes the rates at which molecules interact to protein complexes or dimers that can reversibly dissociate back into component molecules: $T + A \leftrightarrow TA$. Here, T represents thymosin monomers, A

represents actin monomers, TA represents thymosin-actin dimers. The rate equation describing the change in dimer concentration over time Δt is $TA/\Delta t = K_a[T][A] - K_d[TA]$, where values in brackets represent molecular concentrations, K_a represents the forward rate constant, and K_d represents backward rate constant. Cytosolic cycling of thymosin and actin was simulated at a rate: $C = kT/(6\pi\eta r) \times 10^6$, where k is the Boltzmann constant, temperature $T = 310$ kelvin, centipoise viscosity $\eta = 3$, and particle radius $r = 3e-9$ meters

We simulate arp-mediated filament branching using empirically reported branch nucleation rates (Smith et al., 2013). Arp is estimated to nucleate a new filament branch at a rate of $2.5 \text{ mM}_{\text{arp}}^{-1} \mu\text{m}_F^{-1} \Delta t^{-1}$. That is 2.5 new branches per mM free arp, per micron of existing filament, per second. Each new filament branches off the existing filament at a 70-degree angle. Cofilin-mediated filament severing was simulated using the same general rate principles as arp, using empirical rate constants (Halavatyi et al., 2009).

Figure 2.4: Actin dynamics were simulated in dendritic spine volumes of approximately $1e8 \text{ nm}^3$ at full size steady-state. All quantitative data was averaged over 10 independent runs that were intended to simulate 5 hours of real-world time. More runs could easily be generated, but 10 runs allows for easier visualization of the variance between runs (the stochastic nature of the system). Prior to saving quantitative data for using in plots, filament networks were allowed to reach a molecular and morphological steady-state.

Figure 2.5: MCMC methods were used to generate a spatially resolved model of actin filament networks inside dendritic spine volumes. The spatial coordinates of all filaments were collected each small time-step. When any part of a filament passed within 20 nm of the postsynaptic membrane, those line segments were stored, and used as seed-points for SAP clusters. SAP clustering used a modified version of the Shouval algorithm. In our simulations we used a neighbor-dependent algorithm where lattice on-rate probability was a function of actin filament proximity, and lattice off-rate was neighbor dependent.

For comparison, the Shouval cluster model is simulated using a 2-D logical matrix S , where row and column indices are denoted i and j respectively. This matrix represents a 2-D lattice where the

presence or absence of a receptor is indicated by $S_{ij}=1$ or $S_{ij}=0$ respectively. The lattice *off-rate*, μ , is a constant. So the probability of a lattice location being vacated during each time step Δt is $P_{ij}^{off} = S_{ij} \mu \Delta t$. In the Shouval model, lattice *on-rate* is neighbor-dependent, so insertion probability of any given lattice location S_{ij} depends on the occupancy of nearby locations. A 'field' F , a matrix with the exact same dimensions as S , is used to represent the number of neighbors surrounding each lattice location. In Shouval's sample model, any given field location F_{ij} could take an integer value ranging from 0 (neighbors) to 4 (neighbors), which is computed by convolving C , a simple 3-by-3 convolution matrix, with S (though larger or more complex convolution matrices can be used to simulate more elaborate field effects). Shouval also introduces a so called *lattice repulsion* constant L that is subtracted from the convolution output such that $F_{ij} = \text{conv}(S,C) - L$. Concretely if lattice repulsion were set to $L = 2$, and a particular location S_{ij} had 4 neighbors, then the corresponding field $F_{ij} = 4 - 2$. Finally the conditional on-rate probability, is given by: $P_{ij}^{on} = (1 / (1 + \exp(-\beta F_{ij}))) p r \Delta t$. Here β is the slope of the conditional probability function, and both p and r are constants representing the availability of receptors.

The MCMC code detailing the modified version of the cluster model, where lattice on-rate probability was a function of actin filament proximity, and lattice off-rate was neighbor dependent is also available online: github.com/bradmonk/dissertation.

PART II: EXPERIMENTAL TESTS OF KEY MODEL COMPONENTS: Detailed methods and reagents used for primary culture, glutamate uncaging, 2-photon imaging, FRAP, and FLIM/FRET are described in a recent publication from our lab (Dore et al., 2015); most experiments in this chapter use the same general protocols and exactly the same imaging equipment and rig.

Figure 2.6: To prepare primary cultures, the hippocampal regions from P0 Sprague-Dawley rats were dissected, sectioned, filtered, and resuspended in plating media, and plated on PDL-coated glass coverslips in 12-well plates. Every 2-4 days, half the culture media was replaced (Neurobasal-A, FBS,

Pen/Strep, Glutamax). For further description of primary culture methodology see (Nault & Koninck, 2009). Animal procedures were approved by UC San Diego IACUC.

To perform overexpression experiments, pCI-Actin-GFP and pCI-GFP were acquired from addgene.org nonprofit plasmid repository. Plasmid transfection was performed 4 days prior to imaging. Transfection was performed using lipofectamine 2000 reagent according to manufacturer's protocol. Each tissue coverslip was incubated for 4 hours with 2 μg DNA per 4 μL lipofectamine, and then transferred to fresh culture media. Dendrite morphological features were measured and quantified using custom MATLAB software: github.com/bradmonk/neuromorph.

Detailed FRAP methods are described in (Dore et al., 2015). For the FRAP experiment presented in this chapter, images were collected at the following minute-marker relative to photobleaching: -1, 1, 5, 10, 15, 20, 30, 60 minutes. This was done for each spine. Graphed values of 'fraction recovered' are normalized to the target spine's mean fluorescence intensity 1-minute prior to bleaching.

Figure 2.7: We used Sindbis virus to acutely overexpress actin in organotypic hippocampal slice cultures, and then 18-24h later performed electrophysiological recordings to measure mEPSC amplitude and frequency in control neurons and in neurons overexpressing actin. For detailed methods on the Sindbis virus expression vector see (Makino & Malinow, 2009). Hippocampal slice cultures were prepared from PD 6 rat pups (Stoppini et al., 1991), then maintained in culture for 8 days before Sindbis virus (pSinRep5-GFP or pSinRep5-Actin-GFP) infection. Then as noted above, neurons were recorded at within 24 h post-infection. For detailed methods used to record mini-excitatory postsynaptic currents see (Alfonso et al., 2014). In our electrophysiology preparation, TTX and picrotoxin were added to the media. Voltage clamp was held at -60mV; mEPSC output recorded in pA.

Figure 2.8: Actin dissociates from thymosin peptide during LTP. The polymerization rate of actin filaments depends on the available concentration of free actin monomers. Generally actin polymerizes at a rate of 10 actin subunits per μM each second ($10 \mu\text{M}^{-1}\text{s}^{-1}$). So if the cellular concentration of free actin was 10 μM , each filament would gain ~ 100 monomers per second (Kuhn & Pollard, 2005). The depolymerization rate of actin filaments is independent of actin concentration, and generally proceeds at a

rate of 1 subunit per filament per second (1 s^{-1}). So during a 1-second period, filaments are expected to lose 1 monomer each. Using these general rate parameters, the critical concentration for polymerization would be $0.1 \text{ }\mu\text{M}$; at this concentration polymerization and depolymerization are at a steady state.

In spines, the levels of monomeric G-actin and filamentous F-actin are estimated to be $150 \text{ }\mu\text{M}$ and $500 \text{ }\mu\text{M}$, respectively. At $150 \text{ }\mu\text{M}$, G-actin is maintained at levels far above the critical concentration (C_c) for polymerization in neurons ($\sim 0.1 \text{ }\mu\text{M}$). This is accomplished by the sequestering G-actin by cellular thymosin. Thymosin is a short peptide ($\sim 43 \text{ aa}$) known to reversibly bind G-actin (1:1 stoichiometry) at micromolar affinity (Goldschmidt-Clermont et al., 1992; Irobi et al., 2004; Safer et al., 1990). By virtue of thymosin's reversible binding to G-actin, this peptide inhibits a portion of the monomeric actin pool from freely polymerizing. This raises the overall cellular concentration of G-actin well above the monomeric C_c of $0.1 \text{ }\mu\text{M}$, since a large fraction of G-actin will be tied-up in G-actin:thymosin dimers. If indeed prior measurements of spine G-actin levels are accurate ($150 \text{ }\mu\text{M}$), the reserve pool of G-actin:thymosin would be approximately $149.9 \text{ }\mu\text{M}$ ($150 \text{ }\mu\text{M}$ total G-actin, minus $0.1 \text{ }\mu\text{M}$ C_c G-actin steady-state concentration).

Thymosin and actin overexpression experiments were performed in primary cultures (see methods above), while glutamate uncaging and FLIM experiments were performed using organotypic slice cultures. Thymosin-B4 plasmids: GFP-tagged thymosin $\beta 4$ (Tmsb4x-GFP) was ordered from OriGene. To examine thymosin-actin interactions, we replaced the thymosin GFP tag with one of two red fluorescent proteins (RFP): mApple (ex/em λ : 568/592; EC:75k QY:0.49) and mCherry (ex/em λ : 587/610; EC:72k QY:0.22). Two RFPs were tested to assure eGFP (ex/em λ : 488/507; EC:56k QY:0.60) donor compatibility. For transfection in primary culture, products were inserted into pCI mammalian expression vector (Promega, Madison, WI); methods for transfection in primary culture are described above.

Uncaging and FRET/FLIM experiments were performed using slice cultures. Sindbis viral vectors were used to express or coexpress recombinant constructs (pSinRep5-GFP, pSinRep5-Actin-GFP, pSinRep5-Thymosin-mApple) in slice culture (see above for general Sindbis infection protocol). The 1-

photon uncaging of RuBi-Glutamate (30 μ M) was used to stimulate single dendritic spines and induce LTP; immediately after uncaging, detailed images were captured using 2-photon microscopy.

FLIM image acquisition apparatus, software, and methods are described in (Dore et al., 2015), with the following protocol modification: FLIM images were collected for 30 seconds followed by a 30 second gap, starting 60 seconds prior to RuBi-Glu uncaging, then immediately before uncaging, immediately after uncaging, 60 seconds after uncaging. FLIM analysis was performed using a custom analysis library coded in MATLAB: github.com/bradmonk/FLIM . Dendritic spines were selected from a maximum projection image so that analyst is blind to primary dependent variable (lifetime) when selecting spine ROIs. For each spine an adjacent segment of dendrite is also selected for normalization purposes. For each ROI trace, mean lifetime, fluorescence intensity, and spine area are determined.

General & Analysis Methods: Primary components of this model include (1) a 3D mesh representing a segment of dendrite with several spines, (2) surface receptors that can diffuse along the dendrite membrane mesh, (3) a dynamic actin filament network within spine regions, and (4) multivalent SAPs that interact with actin filaments and surface receptors. Surface receptor diffusion was simulated on a 3D dendritic mesh that included several prototypical spines. Diffusion coefficients for surface receptors were set to values reported in single-particle tracking studies that measured AMPAR diffusion along the dendritic shaft, extrasynaptic spine areas, and postsynaptic densities. Receptor diffusion could be reduced synapses through stochastic interactions with SAPs. Clusters of SAPs could form around actin filaments near the postsynaptic membrane using Shouval's clustering algorithm (Shouval, 2005).

All simulated molecules interacted probabilistically based on spatial proximity and empirical K_a/K_d rate parameters; the primary particles (surface receptors, SAP, actin) could perform multivalent interactions: receptors could interact with multiple SAPs; a SAP could interact with a receptor, another SAP, and actin filaments; actin could interact with SAP, other actin monomers, and actin binding proteins. The model was developed using the Python programming language (open-source) and MATLAB programming language and software (MathWorks Inc., Natick, MA, 2011 160); model source-code and additional synthesis details are provided online at github.com/bradmonk/dissertation .

Chapter 2, in part, is currently being prepared for submission for publication of the material. Monk, Bradley; Dore, Kim; Proulx, Christophe; Alphonso, Stephanie; Marino, Marc; Aronson, Sage; Malinow, Roberto. The dissertation author was the primary investigator and author of this material.

WORKS CITED

- Alfonso, S., Kessels, H. W., Banos, C. C., Chan, T. R., Lin, E. T., Kumaravel, G., Scannevin, R. H., Rhodes, K. J., Haganir, R., Guckian, K. M., Dunah, A. W., & Malinow, R. (2014). Synaptic depressive effects of amyloid beta require PICK1. *European Journal of Neuroscience*, *39*(7), 1225–1233. <https://doi.org/10.1111/ejn.12499>
- Bassani, S., Folci, A., Zapata, J., & Passafaro, M. (2013). AMPAR trafficking in synapse maturation and plasticity. *Cellular and Molecular Life Sciences*, *70*(23), 4411–4430. <https://doi.org/10.1007/s00018-013-1309-1>
- Biess, A., Korkotian, E., & Holcman, D. (2007). Diffusion in a dendritic spine: The role of geometry. *Physical Review E*, *76*(2), 021922. <https://doi.org/10.1103/physreve.76.021922>
- Bindschadler, M. (2010). Modeling actin dynamics. *Wiley Interdisciplinary Reviews: Systems Biology and Medicine*, *2*(4), 481–488. <https://doi.org/10.1002/wsbm.62>
- Bindschadler, M., Osborn, E. A., Dewey, C. F., & McGrath, J. L. (2004). A mechanistic model of the actin cycle. *Biophysical Journal*, *86*(5), 2720–2739. [https://doi.org/10.1016/s0006-3495\(04\)74326-x](https://doi.org/10.1016/s0006-3495(04)74326-x)
- Borgdorff, A. J., & Choquet, D. (2002). Regulation of AMPA receptor lateral movements. *Nature*, *417*(6889), 649–653. <https://doi.org/10.1038/nature00780>
- Bredt, D. S., & Nicoll, R. A. (2003). AMPA Receptor Trafficking at Excitatory Synapses. *Neuron*, *40*(2), 1–19. [https://doi.org/10.1016/s0896-6273\(03\)00640-8](https://doi.org/10.1016/s0896-6273(03)00640-8)
- Chen, L. Y., Rex, C. S., Casale, M. S., Gall, C. M., & Lynch, G. (2007). Changes in Synaptic Morphology Accompany Actin Signaling during LTP. *The Journal of Neuroscience*, *27*(20), 5363–5372. <https://doi.org/10.1523/jneurosci.0164-07.2007>
- Cingolani, L. A., & Goda, Y. (2008). Actin in action: the interplay between the actin cytoskeleton and synaptic efficacy. *Nature Reviews Neuroscience*, *9*(5), 344–356. <https://doi.org/10.1038/nrn2373>
- Collingridge, G. L., Isaac, J. T. R., & Wang, Y. T. (2004). Receptor trafficking and synaptic plasticity.

- Nature Reviews Neuroscience, 5(12), 952–962. <https://doi.org/10.1038/nrn1556>
- Crick, F. (1984). Neurobiology: Memory and molecular turnover. *Nature*, 312(5990), 101–101. <https://doi.org/10.1038/312101a0>
- Czöndör, K., Mondin, M., Garcia, M., Heine, M., Frischknecht, R., Choquet, D., Sibarita, J.-B., & Thoumine, O. R. (2012). Unified quantitative model of AMPA receptor trafficking at synapses. *Proceedings of the National Academy of Sciences of the United States of America*, 109(9), 3522–3527. <https://doi.org/10.1073/pnas.1109818109>
- Derkach, V. A., Oh, M. C., Guire, E. S., & Soderling, T. R. (2007). Regulatory mechanisms of AMPA receptors in synaptic plasticity. *Nature Reviews Neuroscience*, 8(2), 101–113. <https://doi.org/10.1038/nrn2055>
- Dore, K., Aow, J., & Malinow, R. (2015). Agonist binding to the NMDA receptor drives movement of its cytoplasmic domain without ion flow. *Proceedings of the National Academy of Sciences*, 112(47), 14705–14710. <https://doi.org/10.1073/pnas.1520023112>
- Ehlers, M. D., Heine, M., Groc, L., Lee, M.-C., & Choquet, D. (2007). Diffusional Trapping of GluR1 AMPA Receptors by Input-Specific Synaptic Activity. *Neuron*, 54(3), 447–460. <https://doi.org/10.1016/j.neuron.2007.04.010>
- Ehrlich, I., Klein, M., Rumpel, S., & Malinow, R. (2007). PSD-95 is required for activity-driven synapse stabilization. *Proceedings of the National Academy of Sciences*, 104(10), 4176–4181. <https://doi.org/10.1073/pnas.0609307104>
- Einstein, A. (1956). *Investigations on the Theory of the Brownian Movement*. Dover.
- Fischer, M., Kaech, S., Knutti, D., & Matus, A. (1998). Rapid Actin-Based Plasticity in Dendritic Spines. *Neuron*, 20(5), 847–854. [https://doi.org/10.1016/s0896-6273\(00\)80467-5](https://doi.org/10.1016/s0896-6273(00)80467-5)
- Fukazawa, Y., Saitoh, Y., Ozawa, F., Ohta, Y., Mizuno, K., & Inokuchi, K. (2003a). Hippocampal LTP Is Accompanied by Enhanced F-Actin Content within the Dendritic Spine that Is Essential for Late LTP Maintenance In Vivo. *Neuron*, 38(3), 447–460. [https://doi.org/10.1016/s0896-6273\(03\)00206-x](https://doi.org/10.1016/s0896-6273(03)00206-x)
- Goldschmidt-Clermont, P. J., Furman, M. I., Wachsstock, D., Safer, D., Nachmias, V. T., & Pollard, T. D. (1992). The control of actin nucleotide exchange by thymosin beta 4 and profilin. A potential regulatory mechanism for actin polymerization in cells. *Molecular Biology of the Cell*, 3(9), 1015–1024. <https://doi.org/10.1091/mbc.3.9.1015>
- Halavaty, A. A., Nazarov, P. V., Medves, S., Troys, M. van, Ampe, C., Yatskou, M., & Friederich, E.

- (2009). An integrative simulation model linking major biochemical reactions of actin-polymerization to structural properties of actin filaments. *Biophysical Chemistry*, 140(1–3), 24–34. <https://doi.org/10.1016/j.bpc.2008.11.006>
- Holcman, D., & Schuss, Z. (2011). Diffusion laws in dendritic spines. *The Journal of Mathematical Neuroscience*, 1(1), 10. <https://doi.org/10.1186/2190-8567-1-10>
- Holcman, D., & Triller, A. (2006). Modeling Synaptic Dynamics Driven by Receptor Lateral Diffusion. *Biophysical Journal*, 91(7), 2405–2415. <https://doi.org/10.1529/biophysj.106.081935>
- Honkura, N., Matsuzaki, M., Noguchi, J., Ellis-Davies, G. C. R., & Kasai, H. (2008). The Subspine Organization of Actin Fibers Regulates the Structure and Plasticity of Dendritic Spines. *Neuron*, 57(5), 719–729. <https://doi.org/10.1016/j.neuron.2008.01.013>
- Hotulainen, P., & Hoogenraad, C. C. (2010). Actin in dendritic spines: connecting dynamics to function. *The Journal of Cell Biology*, 189(4), 619–629. <https://doi.org/10.1083/jcb.201003008>
- Hoze, N., Nair, D., Hosity, E., & Sieben, C. (2012). *Heterogeneity of AMPA receptor trafficking and molecular interactions revealed by superresolution analysis of live cell imaging*. 109. <https://doi.org/10.1073/pnas.1204589109/-/dcsupplemental>
- Irobi, E., Aguda, A. H., Larsson, M., Guerin, C., Yin, H. L., Burtnick, L. D., Blanchoin, L., & Robinson, R. C. (2004). Structural basis of actin sequestration by thymosin-beta4: implications for WH2 proteins. *The EMBO Journal*, 23(18), 3599–3608. <https://doi.org/10.1038/sj.emboj.7600372>
- Jones, R. (2013). Knockout blow for “memory molecule.” *Nature Reviews Neuroscience*, 14(3), 154–155. <https://doi.org/10.1038/nrn3441>
- Kessels, H. W., Kopec, C. D., Klein, M. E., & Malinow, R. (2009). Roles of stargazin and phosphorylation in the control of AMPA receptor subcellular distribution. *Nature Neuroscience*, 12(7), 888–896. <https://doi.org/10.1038/nn.2340>
- Kopec, C. D., Real, E., Kessels, H. W., & Malinow, R. (2007). GluR1 Links Structural and Functional Plasticity at Excitatory Synapses. *Journal of Neuroscience*, 27(50), 13706–13718. <https://doi.org/10.1523/jneurosci.3503-07.2007>
- Korobova, F., & Svitkina, T. (2010). Molecular architecture of synaptic actin cytoskeleton in hippocampal neurons reveals a mechanism of dendritic spine morphogenesis. *Molecular Biology of the Cell*, 21(1), 165–176. <https://doi.org/10.1091/mbc.e09-07-0596>
- Krucker, T., Siggins, G. R., & Halpain, S. (2000). Dynamic actin filaments are required for stable long-term potentiation (LTP) in area CA1 of the hippocampus. *Proceedings of the National Academy*

- of Sciences, 97(12), 6856–6861. <https://doi.org/10.1073/pnas.100139797>
- Kuhn, J. R., & Pollard, T. D. (2005). Real-time measurements of actin filament polymerization by total internal reflection fluorescence microscopy. *Biophysical Journal*, 88(2), 1387–1402. <https://doi.org/10.1529/biophysj.104.047399>
- Kusumi, A., Suzuki, K. G. N., Kasai, R. S., Ritchie, K., & Fujiwara, T. K. (2011). Hierarchical mesoscale domain organization of the plasma membrane. *Trends in Biochemical Sciences*, 36(11), 604–615. <https://doi.org/10.1016/j.tibs.2011.08.001>
- Lang, C., Barco, A., Zablow, L., Kandel, E. R., Siegelbaum, S. A., & Zakharenko, S. S. (2004). Transient expansion of synaptically connected dendritic spines upon induction of hippocampal long-term potentiation. *Proceedings of the National Academy of Sciences*, 101(47), 16665–16670. <https://doi.org/10.1073/pnas.0407581101>
- Liao, D., Hessler, N. A., & Malinow, R. (1995). Activation of postsynaptically silent synapses during pairing-induced LTP in CA1 region of hippocampal slice. *Nature*, 375(6530), 400–404. <https://doi.org/10.1038/375400a0>
- Lin, W.-H., & Webb, D. J. (2009). Actin and Actin-Binding Proteins: Masters of Dendritic Spine Formation, Morphology, and Function. *The Open Neuroscience Journal*, 3, 54–66. <https://doi.org/10.2174/1874082000903020054>
- Lisman, J. (2003). Actin's Actions in LTP-Induced Synapse Growth. *Neuron*, 38(3), 361–362. [https://doi.org/10.1016/s0896-6273\(03\)00257-5](https://doi.org/10.1016/s0896-6273(03)00257-5)
- Lisman, J. E., & Goldring, M. A. (1988). Feasibility of long-term storage of graded information by the Ca²⁺/calmodulin-dependent protein kinase molecules of the postsynaptic density. *Proceedings of the National Academy of Sciences of the United States of America*, 85(14), 5320–5324. <https://doi.org/10.1073/pnas.85.14.5320>
- Lüscher, C., Xia, H., Beattie, E. C., Carroll, R. C., Zastrow, M. von, Malenka, R. C., & Nicoll, R. A. (1999). Role of AMPA Receptor Cycling in Synaptic Transmission and Plasticity. *Neuron*, 24(3), 649–658. [https://doi.org/10.1016/s0896-6273\(00\)81119-8](https://doi.org/10.1016/s0896-6273(00)81119-8)
- Makino, H., & Malinow, R. (2009). AMPA Receptor Incorporation into Synapses during LTP: The Role of Lateral Movement and Exocytosis. *Neuron*, 64(3), 381–390. <https://doi.org/10.1016/j.neuron.2009.08.035>
- Malinow, R., & Malenka, R. C. (2002). AMPA Receptor Trafficking and Synaptic Plasticity. *Annual Review of Neuroscience*, 25(1), 103–126.

<https://doi.org/10.1146/annurev.neuro.25.112701.142758>

- Nair, D., Hosy, E., Petersen, J. D., Constals, A., Giannone, G., Choquet, D., & Sibarita, J. B. (2013). Super-Resolution Imaging Reveals That AMPA Receptors Inside Synapses Are Dynamically Organized in Nanodomains Regulated by PSD95. *Journal of Neuroscience*, 33(32), 13204–13224. <https://doi.org/10.1523/jneurosci.2381-12.2013>
- Nault, F., & Koninck, P. D. (2009). Protocols for Neural Cell Culture, Fourth Edition. *Springer Protocols Handbooks*, 137–159. https://doi.org/10.1007/978-1-60761-292-6_8
- Okamoto, K., Bosch, M., & Hayashi, Y. (2009). The Roles of CaMKII and F-Actin in the Structural Plasticity of Dendritic Spines: A Potential Molecular Identity of a Synaptic Tag? *Physiology*, 24(6), 357–366. <https://doi.org/10.1152/physiol.00029.2009>
- Opazo, P., Sainlos, M., & Choquet, D. (2012). Regulation of AMPA receptor surface diffusion by PSD-95 slots. *Current Opinion in Neurobiology*, 22(3), 453–460. <https://doi.org/10.1016/j.conb.2011.10.010>
- Otmakhov, N., Griffith, L. C., & Lisman, J. E. (1997). Postsynaptic Inhibitors of Calcium/Calmodulin-Dependent Protein Kinase Type II Block Induction But Not Maintenance of Pairing-Induced Long-Term Potentiation. *Journal of Neuroscience*, 17(14), 5357–5365. <https://doi.org/10.1523/jneurosci.17-14-05357.1997>
- Otmakhov, N., Shirke, A. M., & Malinow, R. (1993). Measuring the impact of probabilistic transmission on neuronal output. *Neuron*, 10(6), 1101–1111. [https://doi.org/10.1016/0896-6273\(93\)90058-y](https://doi.org/10.1016/0896-6273(93)90058-y)
- Poncer, J. C., & Malinow, R. (2001). Postsynaptic conversion of silent synapses during LTP affects synaptic gain and transmission dynamics. *Nature Neuroscience*, 4(10), 989–996. <https://doi.org/10.1038/nn719>
- Raastad, M., Storm, J. F., & Andersen, P. (1992). Putative Single Quantum and Single Fibre Excitatory Postsynaptic Currents Show Similar Amplitude Range and Variability in Rat Hippocampal Slices. *European Journal of Neuroscience*, 4(1), 113–117. <https://doi.org/10.1111/j.1460-9568.1992.tb00114.x>
- Ramachandran, B., & Frey, J. U. (2009). Interfering with the Actin Network and Its Effect on Long-Term Potentiation and Synaptic Tagging in Hippocampal CA1 Neurons in Slices In Vitro. *The Journal of Neuroscience*, 29(39), 12167–12173. <https://doi.org/10.1523/jneurosci.2045-09.2009>
- Renner, M., Choquet, D., & Triller, A. (2009). Control of the Postsynaptic Membrane Viscosity. *Journal of Neuroscience*, 29(9), 2926–2937. <https://doi.org/10.1523/jneurosci.4445-08.2009>

- Sacktor, T. C. (2012). Memory maintenance by PKM ζ — an evolutionary perspective. *Molecular Brain*, 5(1), 31. <https://doi.org/10.1186/1756-6606-5-31>
- Safer, D., Golla, R., & Nachmias, V. T. (1990). Isolation of a 5-kilodalton actin-sequestering peptide from human blood platelets. *Proceedings of the National Academy of Sciences*, 87(7), 2536–2540.
- Sainlos, M., Tigaret, C., Poujol, C., Olivier, N. B., Bard, L., Breillat, C., Thiolon, K., Choquet, D., & Imperiali, B. (2010). Biomimetic divalent ligands for the acute disruption of synaptic AMPAR stabilization. *Nature Chemical Biology*, 7(2), 1–11. <https://doi.org/10.1038/nchembio.498>
- Shen, L., Liang, F., Walensky, L. D., & Huganir, R. L. (2000). Regulation of AMPA receptor GluR1 subunit surface expression by a 4. 1N-linked actin cytoskeletal association. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 20(21), 7932–7940.
- Shi, S.-H., Hayashi, Y., Esteban, J. A., & Malinow, R. (2001). Subunit-Specific Rules Governing AMPA Receptor Trafficking to Synapses in Hippocampal Pyramidal Neurons. *CELL*, 105(3), 1–13. [https://doi.org/10.1016/s0092-8674\(01\)00321-x](https://doi.org/10.1016/s0092-8674(01)00321-x)
- Shouval, H. Z. (2005). Clusters of interacting receptors can stabilize synaptic efficacies. *Proceedings of the National Academy of Sciences*, 102(40), 14440–14445. <https://doi.org/10.1073/pnas.0506934102>
- Si, K., & Kandel, E. R. (2016). The Role of Functional Prion-Like Proteins in the Persistence of Memory. *Cold Spring Harbor Perspectives in Biology*, 8(4), a021774. <https://doi.org/10.1101/cshperspect.a021774>
- Simon, C. M., Hepburn, I., Chen, W., & Schutter, E. D. (2013). The role of dendritic spine morphology in the compartmentalization and delivery of surface receptors. *Journal of Computational Neuroscience*, 36(3), 483–497. <https://doi.org/10.1007/s10827-013-0482-4>
- Simoni, A., Griesinger, C. B., & Edwards, F. A. (2003). Development of Rat CA1 Neurones in Acute Versus Organotypic Slices: Role of Experience in Synaptic Morphology and Activity. *The Journal of Physiology*, 550(1), 135–147. <https://doi.org/10.1113/jphysiol.2003.039099>
- Smith, B. A., Daugherty-Clarke, K., Goode, B. L., & Gelles, J. (2013). Pathway of actin filament branch formation by Arp2/3 complex revealed by single-molecule imaging. *Proceedings of the National Academy of Sciences of the United States of America*, 110(4), 1285–1290. <https://doi.org/10.1073/pnas.1211164110>
- Smolen, P., Baxter, D. A., & Byrne, J. H. (2019). How can memories last for days, years, or a lifetime?

- Proposed mechanisms for maintaining synaptic potentiation and memory. *Learning & Memory*, 26(5), 133–150. <https://doi.org/10.1101/lm.049395.119>
- Song, I., & Huganir, R. L. (2002). Regulation of AMPA receptors during synaptic plasticity. *Trends in Neurosciences*, 25(11), 578–588. [https://doi.org/10.1016/s0166-2236\(02\)02270-1](https://doi.org/10.1016/s0166-2236(02)02270-1)
- Stoppini, L., Buchs, P.-A., & Muller, D. (1991). A simple method for organotypic cultures of nervous tissue. *Journal of Neuroscience Methods*, 37(2), 173–182. [https://doi.org/10.1016/0165-0270\(91\)90128-m](https://doi.org/10.1016/0165-0270(91)90128-m)
- Svoboda, K., Tank, D. W., & Denk, W. (1996). Direct measurement of coupling between dendritic spines and shafts. *Science*, 272(5262), 716–719.
- Volk, L. J., Bachman, J. L., Johnson, R., Yu, Y., & Huganir, R. L. (2013). PKM- ζ is not required for hippocampal synaptic plasticity, learning and memory. *Nature*, 493(7432), 420–423. <https://doi.org/10.1038/nature11802>
- Wyllie, D. J. A., Manabe, T., & Nicoll, R. A. (1994). A rise in postsynaptic Ca²⁺ potentiates miniature excitatory postsynaptic currents and AMPA responses in hippocampal neurons. *Neuron*, 12(1), 127–138. [https://doi.org/10.1016/0896-6273\(94\)90158-9](https://doi.org/10.1016/0896-6273(94)90158-9)
- Xue, B., & Robinson, R. C. (2013). Guardians of the actin monomer. *European Journal of Cell Biology*, 92(10–11), 316–332. <https://doi.org/10.1016/j.ejcb.2013.10.012>
- Yarmola, E. G., Dranishnikov, D. A., & Bubb, M. R. (2008). Effect of profilin on actin critical concentration: a theoretical analysis. *Biophysical Journal*, 95(12), 5544–5573. <https://doi.org/10.1529/biophysj.108.134569>
- Yudowski, G. A., Puthenveedu, M. A., Leonoudakis, D., Panicker, S., Thorn, K. S., Beattie, E. C., & Zastrow, M. von. (2007). Real-Time Imaging of Discrete Exocytic Events Mediating Surface Delivery of AMPA Receptors. *The Journal of Neuroscience*, 27(41), 11112–11121. <https://doi.org/10.1523/jneurosci.2465-07.2007>

GENERAL DISCUSSION

In chapter 1, we presented a neural network model that could estimate polygenic risk scores (PRS) for Alzheimer's Disease, with better accuracy than any model reported to-date. Our model represents an advancement over other polygenic risk models, for several reasons. First, other prevailing models only used GWAS summary data to generate risk scores. This can only account for cumulative main effects of variants, not variant-variant interactions. Our model on the other hand was developed using individual genotype, so it could account for variant interactions. Second, our model leverages the power of artificial neural networks, which are ideal for modeling high dimensional data and fitting complex interactions. As a result, the trained model represents to most powerful classifier reported to-date for identifying Alzheimer's polygenic risk.

We also developed a machine learning-based method (netSNP) to identify the importance of individual SNPs in a complex polygenic classifier's decision making process. The netSNP method can be used to investigate the impact of specific SNPs on NN output. We found that netSNP captured the effect of different APOE genotypes on NN output, and also identified hundreds of SNPs with significant neural network weights. netSNP even identified several variants that have more putative impact on AD risk than the well-known APOE genotypes.

Importantly netSNP was well validated. For example the number of protective variants (as identified by netSNP) harbored by an individual correlated significantly with an individual's age of AD diagnosis, and inversely with Braak score. Likewise while the number of netSNP-identified risk variants harbored by an individual was inversely correlated with AD onset and positively correlated with Braak score. Our data suggest the set of variants identified by netSNP are highly predictive of AD onset age and physiological severity. Our findings suggest that netSNP could be an important tool for finding pathophysiologically relevant genes in AD, and potentially as a general method applicable to other conditions.

In chapter 2 we present a model demonstrating that actin could be a central regulator of synaptic weights. This hypothesis is supported by a series of computational and biological experiments illustrating how temporary signals induce persisting changes to synaptic actin networks, and how these filament networks are maintained for periods far outlasting the lifetime of synaptic molecules. In general metastable actin filaments allow synaptic weights to be modified by transient signals and achieve long lasting stability despite total molecular turnover.

The computational modeling experiments in chapter 2 demonstrate that, despite stochastic turnover of surface receptors, synaptic receptor levels can be mediated by synaptic diffusion rates. If a stable diffusion rate differential is maintained in synaptic areas, a predictable number of surface receptors will accumulate in synapses. To explain how diffusion rate gradients are stable despite constant protein turnover we included actin in our model after identifying its promising metastable properties. Our computational modeling experiments incorporating actin show that actin filament networks can maintain average filament lengths despite the complete turnover of their F-actin subunits. Simulation experiments also show filament lengths can stably increase following transient growth signals.

Biological experiments in chapter 2 also support essential elements of an ‘actin memory’ model. We find actin filament networks in dendritic spines are morphologically stable orders of magnitude longer than individual actin subunits: all spine actin content is replaced in about one hour without spines undergoing significant changes in size. We also find that global overexpression of actin was sufficient to alter synaptic morphology, increase spine size, and increase spine density. Furthermore, upregulating actin bioavailability was sufficient to increase synaptic weights. We found acute actin overexpression in hippocampal neurons significantly increased mEPSC amplitudes compared to control neurons. Lastly, we identify a biological mechanism that produced acute increases in actin bioavailability at single dendritic spines. Thymosin, which sequesters G-actin under basal conditions, was found to release actin in large quantities in response to LTP signals, leading to significant and long-lasting increase in dendritic spine size. Overall, our model involving actin, SAP, and surface receptors, provides an example of stochastic systems within dendrites giving rise to stability in the aggregate. We show through a series of

computational and biological experiments that actin has metastable properties, which allow synaptic weights to be modified by transient signals and attain long-term stability despite total molecular turnover.