



Genetic and Evolutionary Computation Conference



San José, Costa Rica
July 13-17, 2026

Evolution With Purpose: Hierarchy-Informed Optimization of Whole-Brain Models

Hormoz Shahrzad, Niharika Gajawelli, Kaitlin Maile, Manish Saggar, Risto Miikkulainen

Cognizant AI Lab, University of Texas at Austin, Stanford University


San Francisco, CA, USA

hormoz@cs.utexas.edu



What is a Whole-Brain Model?

Input: Structural Connectivity (SC) from Diffusion-Weighted Imaging (DWI) → Dynamic Mean Field (DMF) Neural Simulation →
 Synthetic Blood-Oxygen-Level-Dependent (BOLD) Signal → Functional Connectivity (FC) → Minimizing FC error



Anatomical wiring
 from diffusion MRI (DWI)
 — not a time series —

DMF Model


400 brain regions
 140 free parameters

$$\tau_E \dot{S}_i = -S_i + (1 - S_i) \gamma H(I_i)$$

*Coupled
 excitatory-inhibitory
 neural masses per region*

Evolutionary search
 optimizes parameters
 to minimize FC error

Empirical BOLD Signal




Empirical FC Matrix



↕ Fitness: Pearson r between the FCs ↕

Simulated BOLD Signal



Simulated FC Matrix



Key Points

INPUT:

Anatomical wiring (SC)
 from diffusion MRI — not
 time series

MODEL OUTPUT:

Synthetic BOLD signals
 whose statistics should
 match real fMRI

FITNESS:

Pearson correlation
 between simulated and
 empirical FC

NOT:

Time-series prediction.
 Statistical structure must
 match, not exact
 waveforms

RESULT:


Good-fit parameters
 encode the subject's
 unique brain activity
 signature

SC → FC: Honey et al. (2009); DMF model: Deco et al. (2013), Wong & Wang (2006); Large-scale brain dynamics: Breakspear (2017)

The Challenge: High-Dimensional DMF Model Fitting


The Goal: Biological Realism

Large-scale Dynamic Mean Field (DMF) models require hundreds of interacting parameters to accurately simulate complex resting-state network interactions and reproduce empirical fMRI data.



The Problem: The Curse of Dimensionality

- High non-convexity leads to degeneracy.
- Models fit to highly idiosyncratic solutions.
- Models fail to predict cognitive behavior.



Adapted from https://github.com/ThomasYeoLab/CBIG/blob/master/stable_projects/brain_parcellation/Schaefer2018_LocalGlobal/readme_figures/Schaefer2018_400parcel_parcellation_match_Yeo_7_network_fslr32k.png

The Paradox of Whole-Brain Modeling

Good In-Sample fits for individual subjects



$$\tau_E \frac{dS_i}{dt} = -S_i + (1 - S_i) \gamma H(I_i(t)), \quad (1)$$


$$I_i(t) = I_b + w_E S_i(t) - w_I S_i^I(t) + G \sum_j SC_{ij} S_j(t - d_{ij}), \quad (2)$$

Cross-Subject Generalization



- These solutions are brittle
- Fail to generalize cross subject
- Have no prediction power for human behavior

Guiding Evolutionary Search with Biological Logic



The Macroscale Gradient


The human cortex is organized along a principal hierarchical gradient, ranging from foundational sensory-motor processing to higher-order associative cognition.

The Biological Hypothesis

If the brain is structured hierarchically, optimization should be too. Parameters governing foundational global states must be locked in before tuning higher-order associative systems.

Hierarchy-Informed Curriculum Optimization (HICO)

A domain-derived temporal decomposition: we restrict when parameters can vary without changing model size.



Key Insight: Instead of exploring all 140 parameters simultaneously, HICO builds dynamic stability sequentially.

Five distinct evolutionary strategies were tested to isolate the impact of curriculum


Strategy	Parameters	Optimization Style	Biological Grounding	
Homogeneous (Baseline)	20 shared	Flat optimization. Assumes uniform brain dynamics.	✗	None
Heterogeneous	140 total	Flat optimization. Maximizes search dimensionality immediately.	✗	None
HICO	140 total	Phased transmodal-to-unimodal. Prevents interference.		Biologically informed
Reverse-Phased	140 total	Phased unimodal-to-transmodal optimization. Tests directionality.		Biologically informed
Shuffled	140 total	Randomly phased optimization. Tests structure without alignment.	✗	Random

Note: All strategies were given an identical computational budget (120 generations, population of 100) using an elitist genetic algorithm.

High Dimensionality Helps, But Bad Structure Wastes It

High Dimensionality Helps, But Bad Structure Wastes It

Dimensionality Helps
Moving from 20 shared parameters (Homogeneous) to 140 parameters yields highly significant improvements ($p \approx 10^{-25}$) in matching fMRI data.



Multiple Paths Fit Well
Heterogeneous, HICO, and Reverse-Phased models reach similarly high fitness.


The Trap
Shuffled curricula waste the dimensional advantage.
Random Order Breaks Search (Order matters).

Curriculum Learning Prevents Dynamical Collapse

The Generalization Collapse Under Flat Search

The Leave-One-Out (LOO) Test:

Averaging fitted parameters across a cohort and testing on a new, unseen subject.



The Collapse Zone


Flat search collapses into unstable regimes where the Lyapunov equation is unsolvable.

The Stability Zone

Curriculum-guided search preserves solvable, coherent, high-fitness dynamics.

What Do We Expect From a Good Model?

A truly generalizable model must encode meaningful differences in human cognition



Statistical Validation of HICO Models

Behavioral Signal Survives 10,000 Permutations, FDR-Corrected



The Methodology:

Empirical nulls from 10,000 label permutations.



The Proof:

HICO's R^2 lies in the extreme permutation tail with FDR-corrected significance.




Biological Localization:

Predictive power localizes to biologically relevant cortical networks.



PMAT24_A_CR (Fluid Reasoning) Null Distributions



Some Fits Are More Equal Than Others

Fitness Alone Doesn't Tell the Whole Story

- **The Good**

High-fitness solutions that demonstrate strong behavioral prediction power.


- **The Bad**

Low-fitness solutions that provide no meaningful behavioral prediction power.

- **The Ugly**


High-fitness solutions that appear promising during optimization but lack behavioral prediction power.

PMAT24_A_CR (Fluid Reasoning) Null Distributions



What Comes Next?

Future Works: Scaling HICO Toward Clinical Impact





Genetic and Evolutionary
Computation Conference



San José, Costa Rica

July 13-17, 2026

Thank you!

Evolution With Purpose: Hierarchy-Informed Optimization of Whole-Brain Models

Hormoz Shahrzad, Niharika Gajawelli, Kaitlin Maile, Manish Saggar, Risto Miikkulainen

Cognizant AI Lab, University of Texas at Austin, Stanford University

San Francisco, CA, USA

hormoz@cs.utexas.edu



sigevo



References I



Pablo Achard and Erik De Schutter.

Complex parameter landscape for a complex neuron model.
PLoS computational biology, 2(7):e94, 2006.



Thomas M. Achenbach, Masha Y. Ivanova, and Leslie A. Rescorla.

Empirically based assessment and taxonomy of psychopathology for ages 1 $\frac{1}{2}$ –90+ years: Developmental, multi-informant, and multicultural findings.
Comprehensive Psychiatry, 79:4–18, 2017.
Advances in Transdiagnostic Psychopathology Research.



Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston.

Curriculum learning.
In *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML '09, page 41–48, New York, NY, USA, 2009. Association for Computing Machinery.



Michael Breakspear.

Dynamic models of large-scale brain activity.
Nature Neuroscience, 20(3):340–352, 2017.



Daniel T. Brocker, Braden D. Swan, Dennis A. Turner, Robert E. Gross, and Warren M. Grill.

Optimized temporal pattern of brain stimulation designed by computational evolution.
Science Translational Medicine, 9(371), 2017.



Joana Cabral, Morten L. Kringelbach, and Gustavo Deco.

Exploring the network dynamics underlying brain activity during rest.
Progress in Neurobiology, 114:102–131, 2014.

References II



Kalyanmoy Deb and Christie Myburgh.

A population-based fast algorithm for a billion-dimensional resource allocation problem with integer variables.

European Journal of Operational Research, 261(2):460–474, 2017.



Gustavo Deco, Viktor Jirsa, Anthony R McIntosh, Olaf Sporns, and Rolf Kötter.

Key role of coupling, delay, and noise in resting brain fluctuations.

Proceedings of the National Academy of Sciences, 106(25):10302–10307, 2009.



Gustavo Deco, Viktor K. Jirsa, and Anthony R. McIntosh.

Emerging concepts for the dynamical organization of resting-state activity in the brain.

Nature Reviews Neuroscience, 12(1):43–56, 2011.



Gustavo Deco, Adrián Ponce-Alvarez, Dante Mantini, Gian Luca Romani, Patric Hagmann, and Maurizio Corbetta.

Resting-state functional connectivity emerges from structurally and dynamically shaped slow linear fluctuations.

Journal of Neuroscience, 33(27):11239–11252, 2013.



Murat Demirtaş, Joshua B. Burt, Markus Helmer, Jie Lisa Ji, Brendan D. Adkinson, Matthew F. Glasser, David C. Van Essen, Stamatis N. Sotiropoulos, Alan Anticevic, and John D. Murray.

Hierarchical heterogeneity across human cortex shapes large-scale neural dynamics.

Neuron, 101(6):1181–1194.e13, 2019.

References III



Shaul Druckmann, Yochai Banitt, Albert Gidon, Felix Schürmann, Henry Markram, and Idan Segev.

A novel multiple objective optimization framework for constraining conductance-based neuron models by experimental data.

Frontiers in Neuroscience, 1(1):7–18, 2007.



Karl J. Friston, Lee Harrison, and William Penny.

Dynamic causal modelling.

NeuroImage, 19(4):1273–1302, 2003.



Matthew F. Glasser, Timothy S. Coalson, Emma C. Robinson, Carl D. Hacker, John Harwell, Essa Yacoub, Kamil Ugurbil, Jesper Andersson, Christian F. Beckmann, Mark Jenkinson, Stephen M. Smith, and David C. Van Essen.

A multi-modal parcellation of human cerebral cortex.

Nature, 536(7615):171–178, 2016.



Faustino Gomez and Risto Miikkulainen.

Incremental evolution of complex general behavior.

Adaptive Behavior, 5(3-4):317–342, 1997.



Nikolaus Hansen.

The cma evolution strategy: A tutorial.

arXiv preprint arXiv:1604.00772, 2016.



Nikolaus Hansen and Andreas Ostermeier.

Completely derandomized self-adaptation in evolution strategies.

Evolutionary Computation, 9(2):159–195, 2001.

References IV



Desmond J. Higham.

An algorithmic introduction to numerical simulation of stochastic differential equations.
SIAM Review, 43(3):525–546, 2001.



John H Holland.

Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence.
MIT press, 1992.



Christopher J. Honey, Olaf Sporns, Leila Cammoun, Xavier Gigandet, Jean-Philippe Thiran, Reto Meuli, and Patric Hagmann.

Predicting human resting-state functional connectivity from structural connectivity.
Proceedings of the National Academy of Sciences, 106(6):2035–2040, 2009.



Peter J. Huber and Elvezio M. Ronchetti.

Robust Statistics.
Wiley Series in Probability and Statistics. John Wiley & Sons, Hoboken, NJ, 2nd edition, 2009.



Naomi Keren, Noam Peled, and Alon Korngreen.

Automated parameter estimation of multi-compartment models of single neurons.
Journal of Neuroscience Methods, 146:133–144, 2005.



Xiaodong Ma, Yuhui Li, Dunwei Gong, Jun Zhang, and Jun Liu.

A survey on cooperative co-evolutionary algorithms.
IEEE Transactions on Evolutionary Computation, 23(3):421–441, 2019.

References V



Kaitlin Maile, Manish Saggar, and Risto P. Miikkulainen.

Implementing evolutionary optimization to model neural functional connectivity.

In *Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO '19 Companion)*, pages 1731–1733. ACM, 2019.



Daniel S. Margulies, Satrajit S. Ghosh, Alexandros Goulas, Magdalena Falkiewicz, Julia M. Huntenburg, Georg Langs, Gleb Bezgin, Simon B. Eickhoff, F. Xavier Castellanos, Michael Petrides, Elizabeth Jefferies, and Jonathan Smallwood.

Situating the default-mode network along a principal gradient of macroscale cortical organization.

Proceedings of the National Academy of Sciences, 113(44):12574–12579, 2016.



Leland McInnes, John Healy, Nathaniel Saul, and Lukas Großberger.

Umap: Uniform manifold approximation and projection.

Journal of Open Source Software, 3(29):861, 2018.



Risto Miikkulainen.

Creative ai through evolutionary computation.

In *Evolution in Action: Past, Present and Future: A Festschrift in Honor of Erik D. Goodman*, pages 265–269. Springer, 2020.



Tyler M Moore, Steven P Reise, Raquel E Gur, Hakon Hakonarson, and Ruben C Gur.

Psychometric properties of the penn computerized neurocognitive battery.

Neuropsychology, 29(2):235, 2015.

References VI



John D Murray, Alberto Bernacchia, David J Freedman, Ranulfo Romo, Jonathan D Wallis, Xinying Cai, Camillo Padoa-Schioppa, Tatiana Pasternak, Hyojung Seo, Daeyeol Lee, et al.
A hierarchy of intrinsic timescales across primate cortex.
Nature neuroscience, 17(12):1661–1663, 2014.



Martin Pelikan.
Bayesian Optimization Algorithm, pages 31–48.
Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.



Mitchell A. Potter and Kenneth A. De Jong.
Cooperative coevolution: An architecture for evolving coadapted subcomponents.
Evolutionary Computation, 8(1):1–29, 2000.



Alexander Schaefer, Ru Kong, Evan M Gordon, Timothy O Laumann, Xi-Nian Zuo, Avram J Holmes, Simon B Eickhoff, and B T Thomas Yeo.
Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri.
Cerebral Cortex, 28(9):3095–3114, 07 2018.



Hormoz Shahrzad, Daniel Fink, and Risto Miikkulainen.
Enhanced optimization with composite objectives and novelty selection.
In *Artificial Life Conference Proceedings*, pages 616–622. MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA, 07 2018.

References VII



Hormoz Shahrzad, Babak Hodjat, Camille Dollé, Andrei Denissov, Simon Lau, Donn Goodhew, Justin Dyer, and Risto Miikkulainen.

Enhanced optimization with composite objectives and novelty pulsation.

In Proceedings of the Genetic and Evolutionary Computation Conference (GECCO), pages 507–515, 2019.



Hormoz Shahrzad and Risto P Miikkulainen.

Accelerating evolution through gene masking and distributed search.

In Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '23, page 972–980, New York, NY, USA, 2023. Association for Computing Machinery.



Kenneth O. Stanley and Risto Miikkulainen.

Evolving neural networks through augmenting topologies.

Evolutionary Computation, 10(2):99–127, 2002.



Werner Van Geit, Michael Gevaert, Giuseppe Chindemi, Christian Rössert, Jean-Denis Courcol, Eilif B. Muller, Felix Schürmann, Idan Segev, and Henry Markram.

Bluepyopt: Leveraging open source software and cloud infrastructure to optimize model parameters in neuroscience.

Frontiers in Neuroinformatics, 10:17, 2016.



Marie C. Vanier and James M. Bower.

A comparative survey of automated parameter-search methods for compartmental neural models.

Journal of Computational Neuroscience, 7(2):149–171, 1999.

References VIII



Miles Wischnewski, Fikret M. Atay, Viktor K. Jirsa, and Gustavo Deco.
Evolutionary optimization of whole-brain models.
Patterns, 3(9), 2022.



Kong-Fatt Wong and Xiao-Jing Wang.
A recurrent network mechanism of time integration in perceptual decisions.
Journal of Neuroscience, 26(4):1314–1328, 2006.



B. T. Thomas Yeo, Fenna M. Krienen, Jorge Sepulcre, Mert R. Sabuncu, Danial Lashkari, Marisa Hollinshead, Joshua L. Roffman, Jordan W. Smoller, Lilla Zöllei, Jonathan R. Polimeni, Bruce Fischl, Hesheng Liu, and Randy L. Buckner.
The organization of the human cerebral cortex estimated by intrinsic functional connectivity.
Journal of Neurophysiology, 106(3):1125–1165, 2011.