

While technological revolutions in neuroscience now enable us to record from ever increasing numbers of neurons, the number we will be able to record in the foreseeable future remains an infinitesimal fraction of the total number of neurons in mammalian circuits controlling complex behaviors. Nevertheless, despite operating within this extreme under-sampling limit, a wide array of statistical procedures for dimensionality reduction of multineuronal recordings uncover remarkably insightful, low dimensional neural state space dynamics whose geometry reveals how behavior and cognition emerge from neural circuits. What theoretical principles explain this remarkable success; in essence, how is it that we can understand anything about the brain while recording an infinitesimal fraction of its degrees of freedom?

We develop an experimentally testable theoretical framework to answer this question. By making a novel conceptual connection between neural measurement and the theory of random projections, we derive scaling laws relating how many neurons we must record to accurately recover state space dynamics, given the complexity of the behavioral or cognitive task, and the smoothness of neural dynamics. Moreover we verify these scaling laws in the motor cortical dynamics of monkeys performing a reaching task.

Along the way, we derive new upper bounds on the number of random projections required to preserve the geometry of smooth random manifolds to a given level of accuracy. Our methods combine probability theory with Riemannian geometry to improve upon previously described upper bounds by two orders of magnitude.