Let the Physics Do the Work: Scattering Algorithms for High-Dimensional Geometry

Aaron Denney, UNM
Cristopher Moore, UNM / Santa Fe Institute

What is the algorithm (what do we design)? What is the input?

Quantum circuits

- algorithm: a series of unitary operators
- input: a unitary operator that queries an oracle

Adiabatic quantum computing

- algorithm: an initial Hamiltonian and an interpolation scheme
- input: a term in the final Hamiltonian

Quantum walks, a.k.a. scattering algorithms

- input: a Hamiltonian, e.g. the adjacency matrix of a graph
- algorithm: Schrödinger's equation

Continuous time: Analog Analogue of Grover [Farhi & Gutmann 98]

An unknown marked state *v*, and a known initial state *u*

A Hamiltonian

$$H = |u\rangle\langle u| + |v\rangle\langle v|$$

The algorithm:

$$-i\hbar \frac{\partial}{\partial t}\psi = H\psi\,.$$

Start in *u*. At time

$$t = \frac{\pi}{2\,\langle u, v\rangle}$$

the state is *v*!

Quantum walks

Input: a Hamiltonian, e.g., the adjacency matrix of a graph

Algorithm: Schrödinger's equation

- Mixing on the cycle [Ambainis et al., Aharonov et al.]
- Mixing and hitting on the hypercube [Moore, Russell; Kempe]
- Exponential speedup possible [Childs et al.]
- Element distinctness [Ambainis]
- Search in d dimensions [Ambainis, Kempe, Rivosh; Childs, Goldstone]
- Hidden nonlinear structures [Childs, Schulman, Vazirani]
- NAND trees [Farhi, Goldstone, Gutmann]
- Boolean formulas [Reichardt, Spalek; Ambainis et al.]

Quantum walks

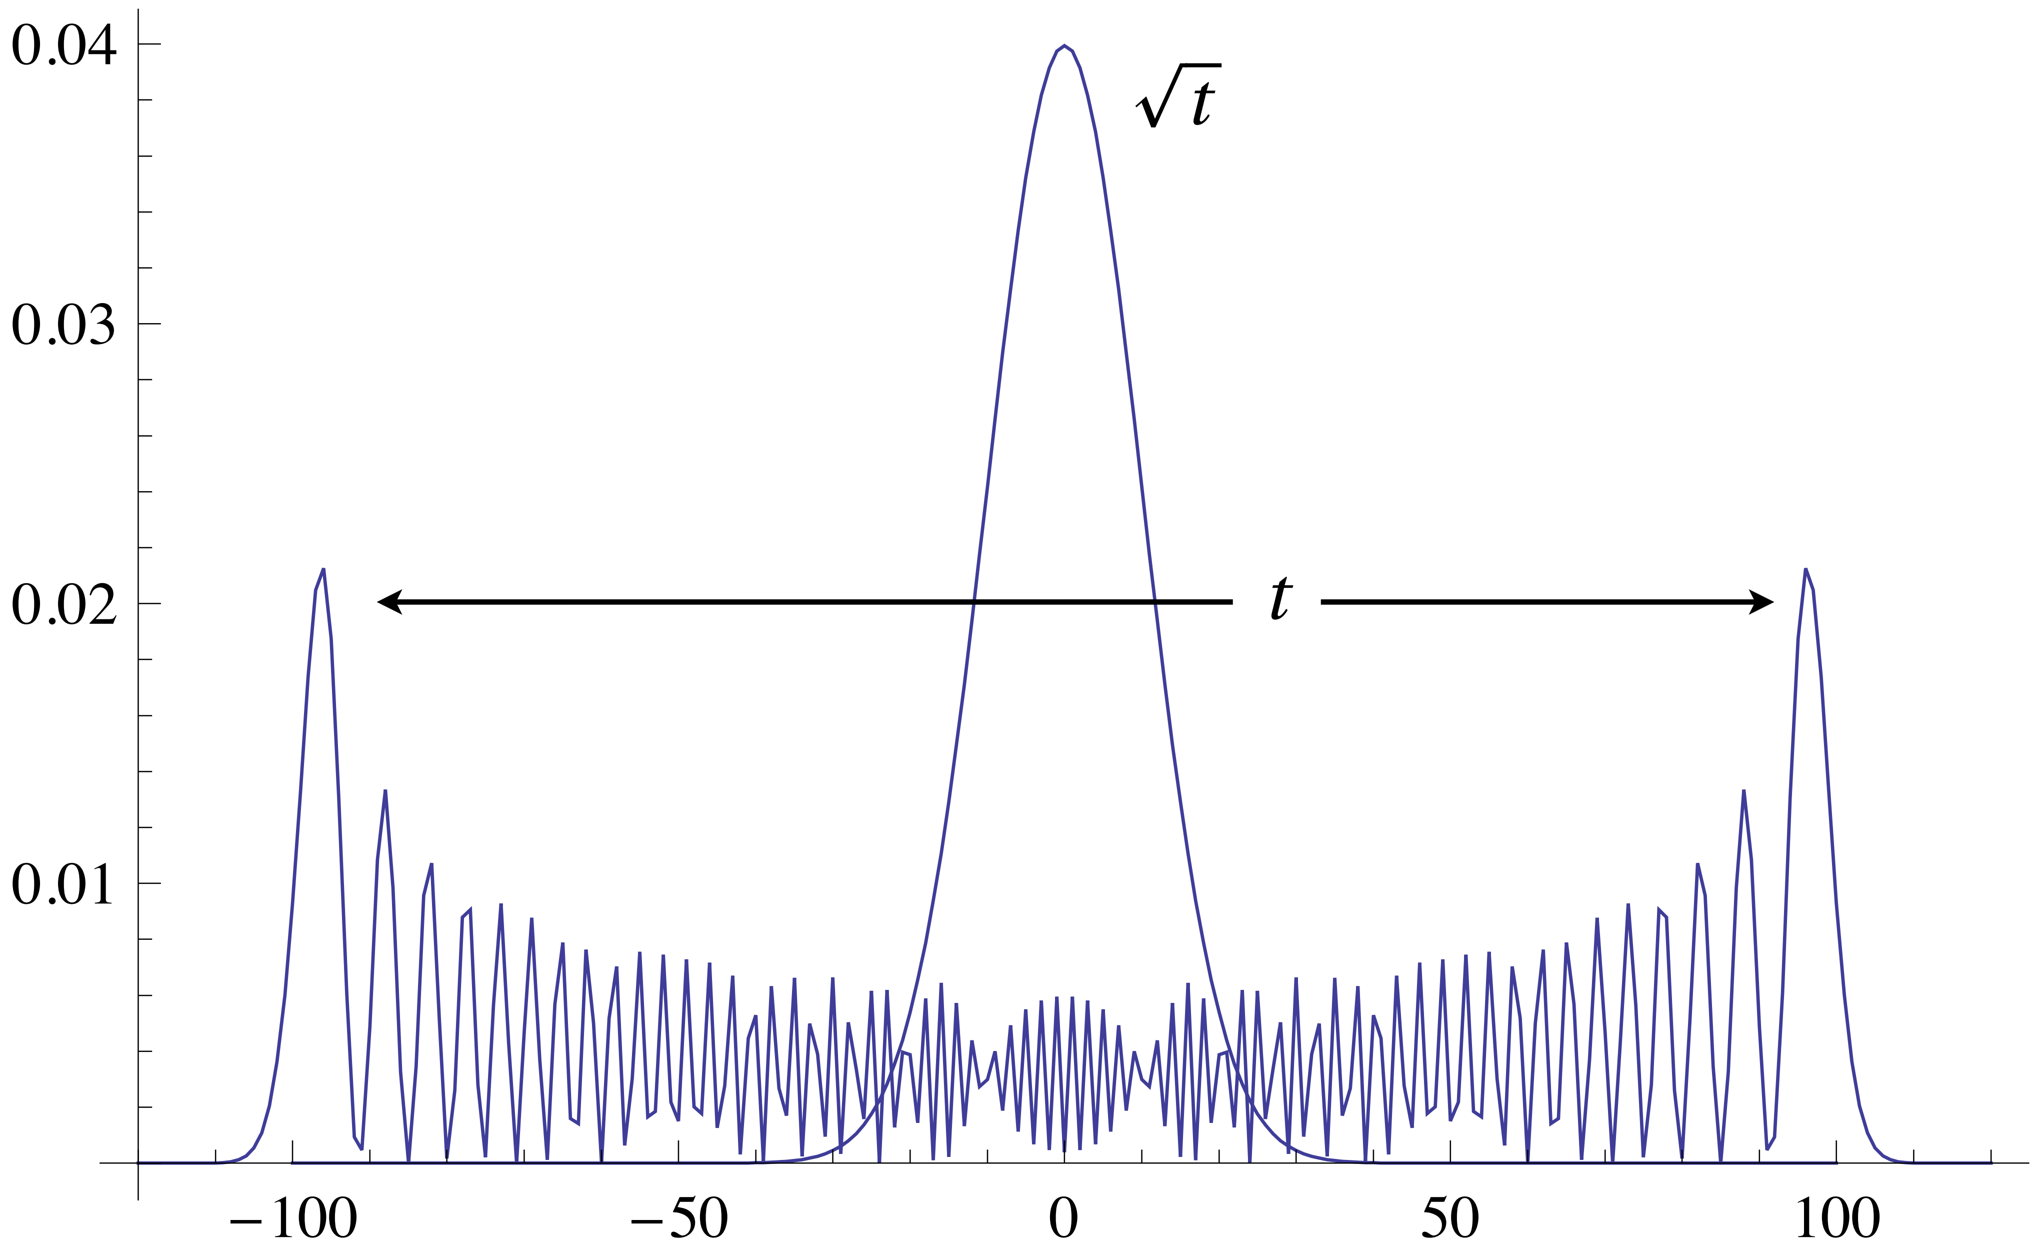

Quantum walks

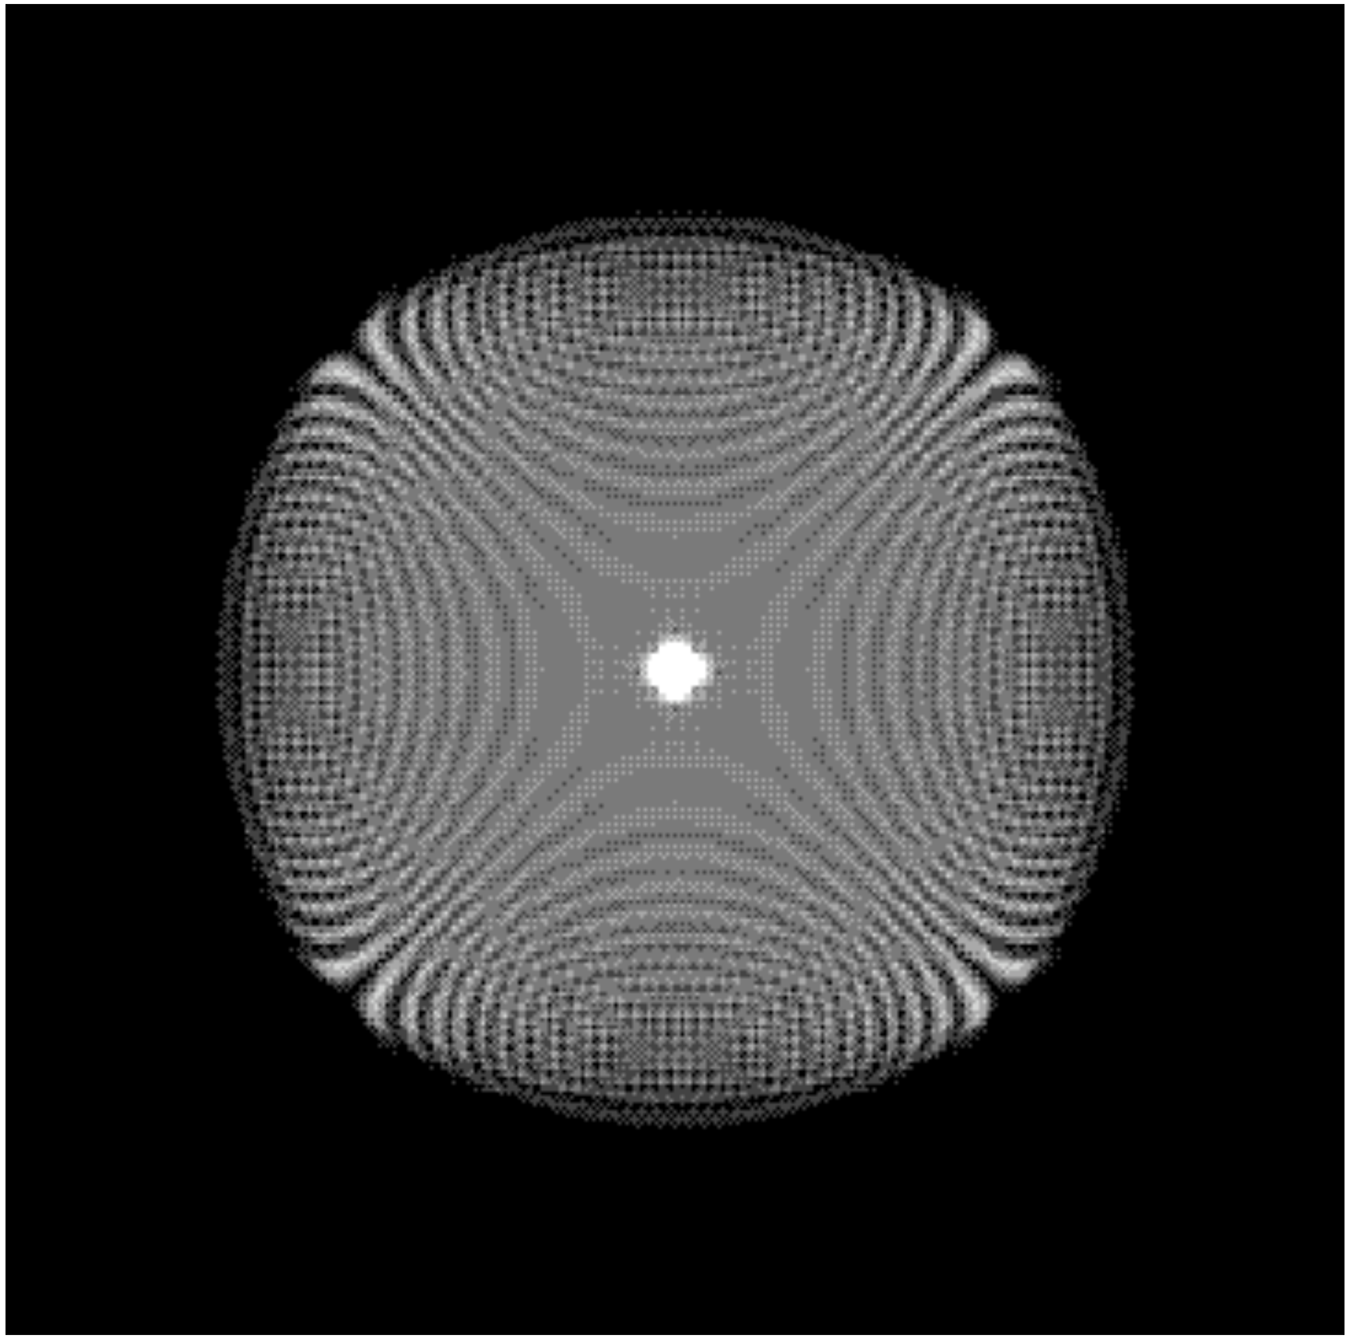

NAND trees in motion

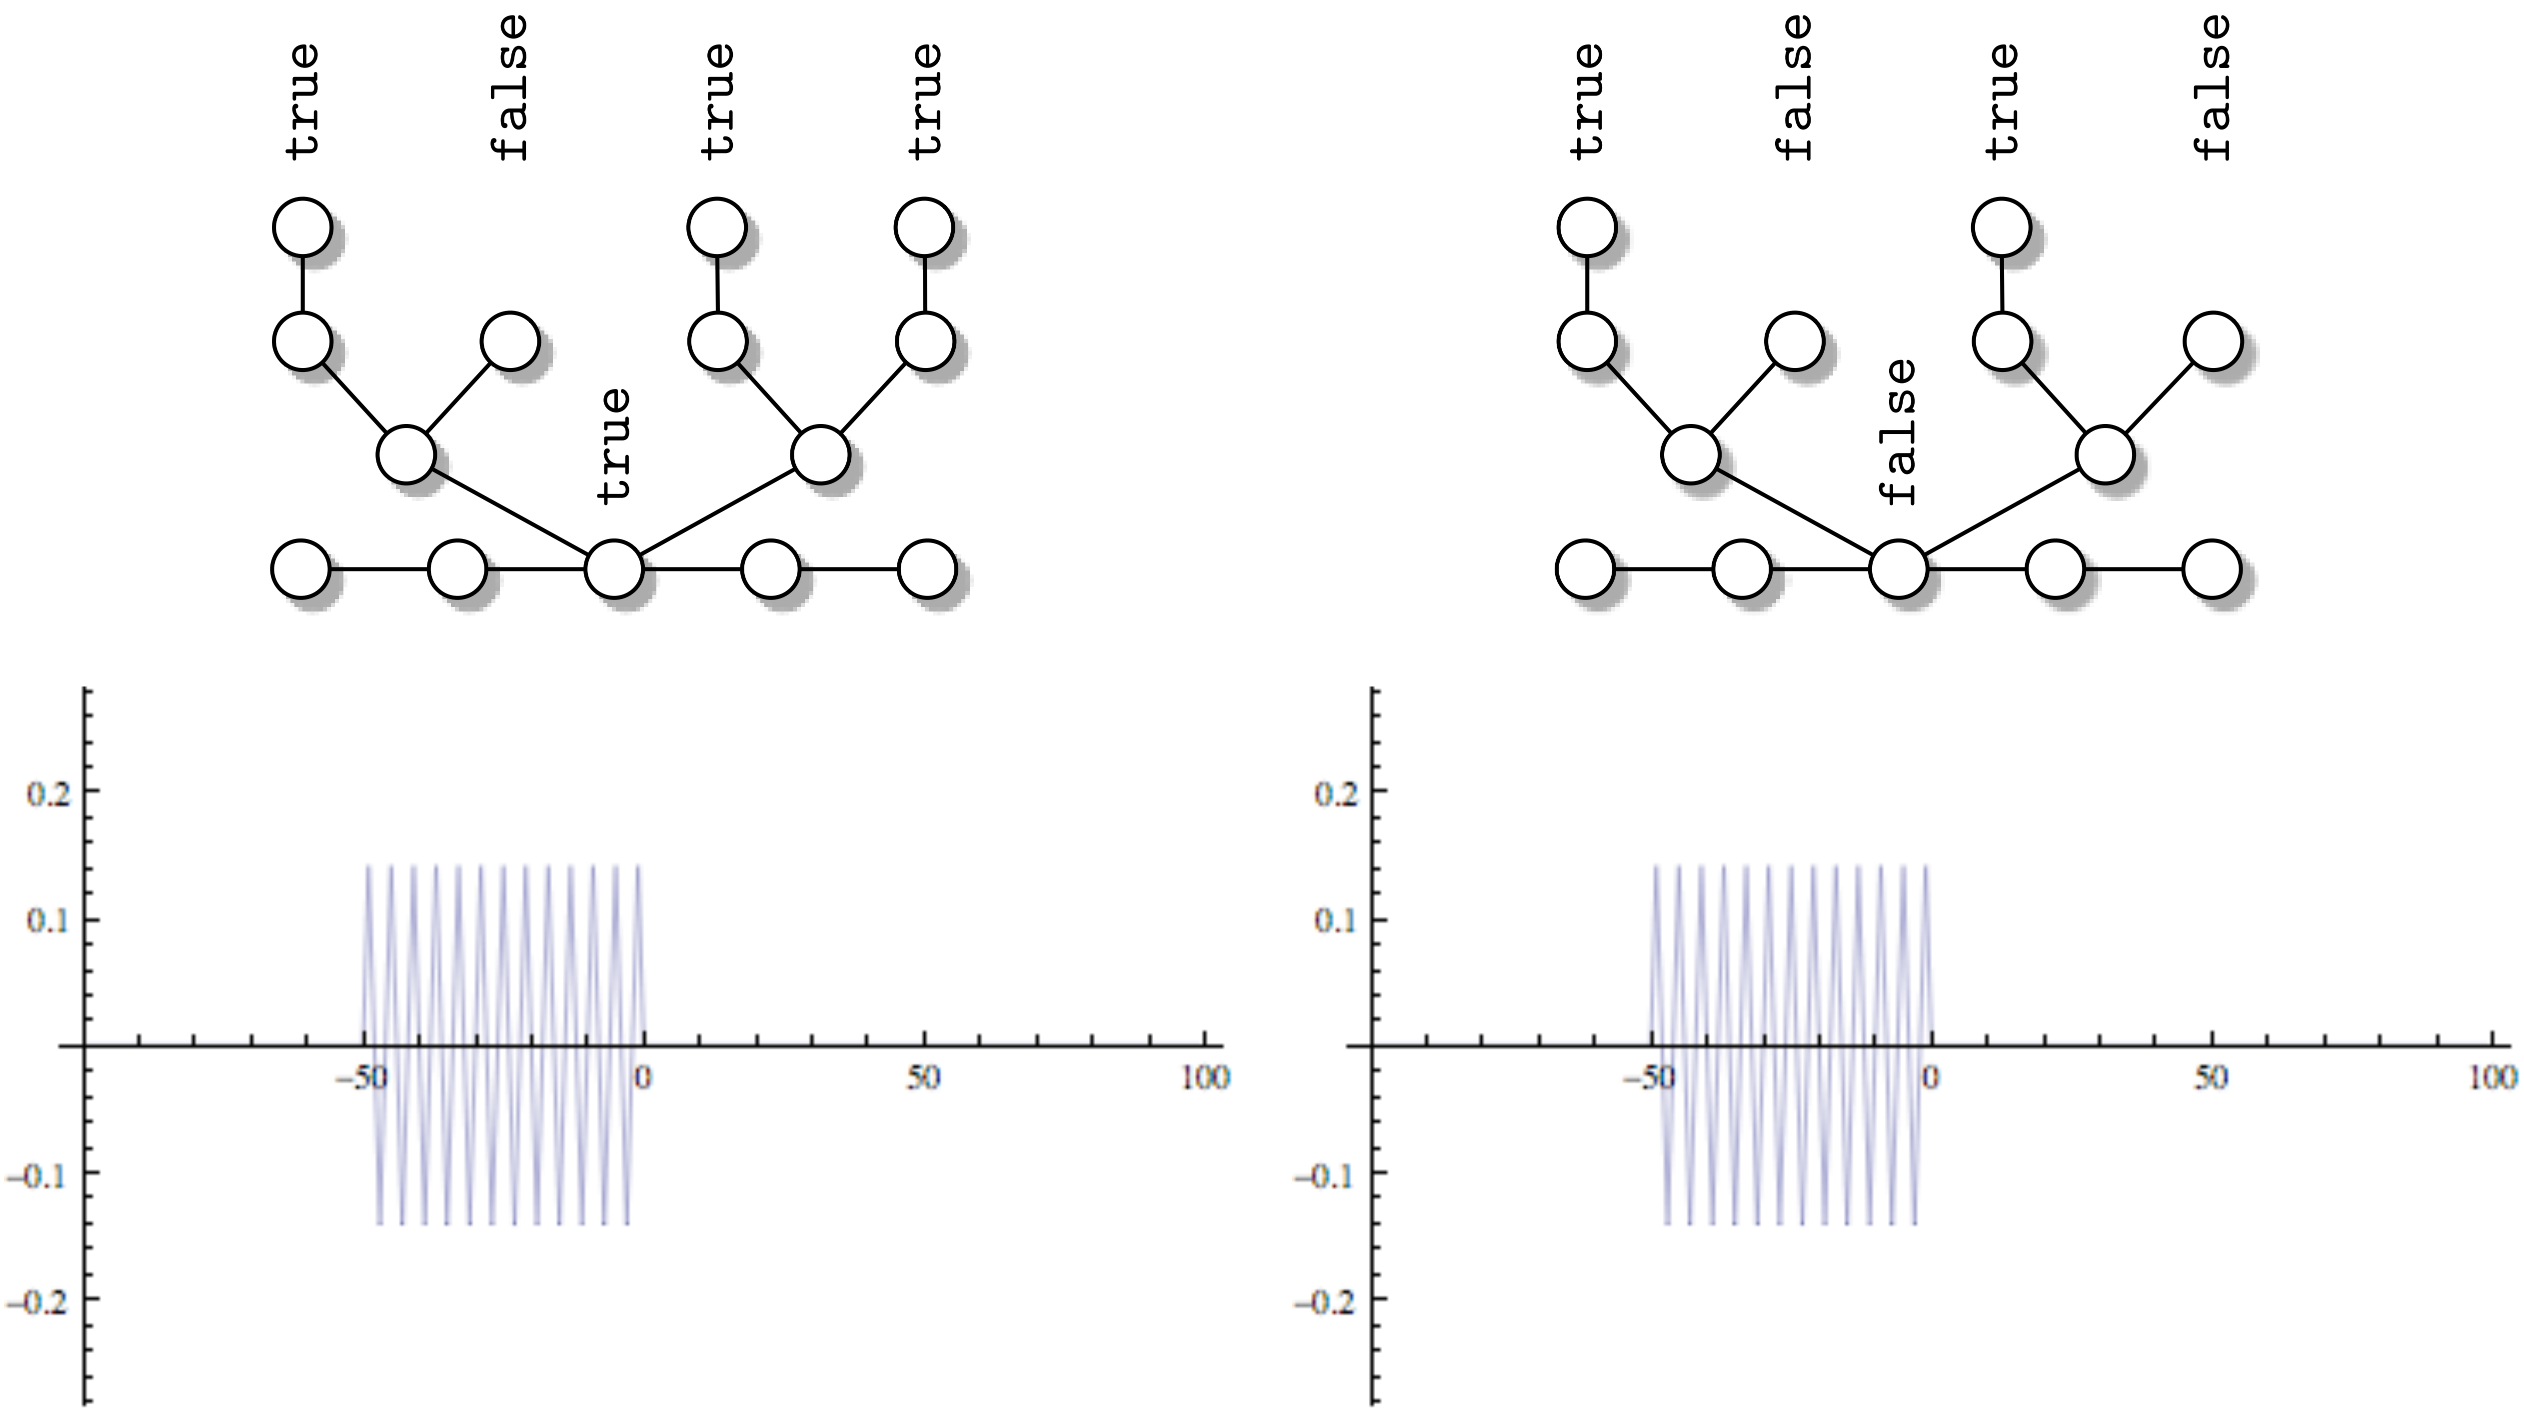

First-semester quantum algorithms

Hamiltonian for a free particle in R^n with zero potential:

$$H = \frac{p^2}{2m} = -\hbar^2 \frac{\nabla^2}{2m} .$$

Equivalent to quantum walk on the *n*-dimensional lattice in the low-energy limit

What can we learn about the initial state, simply by running Schrödinger's equation for a certain amount of time and measuring the position?

Spherical shells

Suppose the initial state is concentrated on a spherical shell of unknown center and/or unknown radius

Application: finding the center of symmetry of a spherically symmetric function in R^n

Let the initial thickness of the shell be $w_0 \sim 1/\tau$

Evolve for time $t \sim w_0^2\tau$, so the width of the state is roughly the radius of the sphere

As in geometric diffraction [Fresnel, Poisson, Arago] we get a bright spot at the center

Independent of the radius (up to a constant)

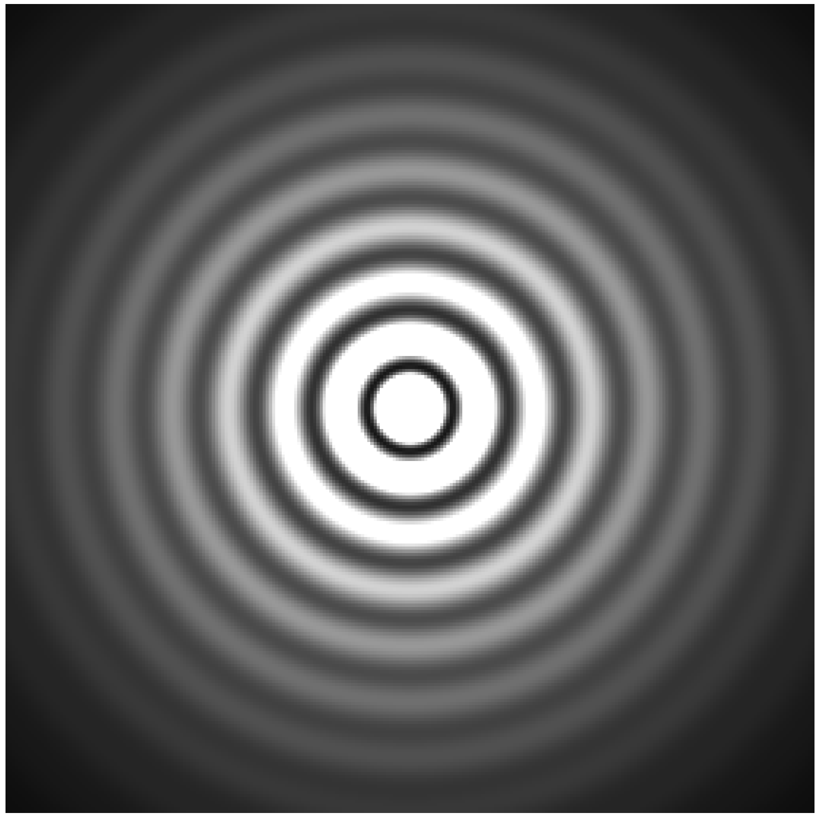

How bright? A singularity at the center

Density *r* away from the center: if $r \gg nw_0$ the probability density scales as

$$|\psi|^2 \sim r^{-(n-1)}$$

so the probability density of the distance *r* is independent of *n*:

$$\rho(r) \sim 1$$

the total probability we observe a point at $r < \varepsilon$ is

$$\Pr[r < \varepsilon] \sim \varepsilon$$

if $|\psi|^2$ were smooth, this would be exponentially small

$$\Pr[r < \varepsilon] \sim \varepsilon^n$$

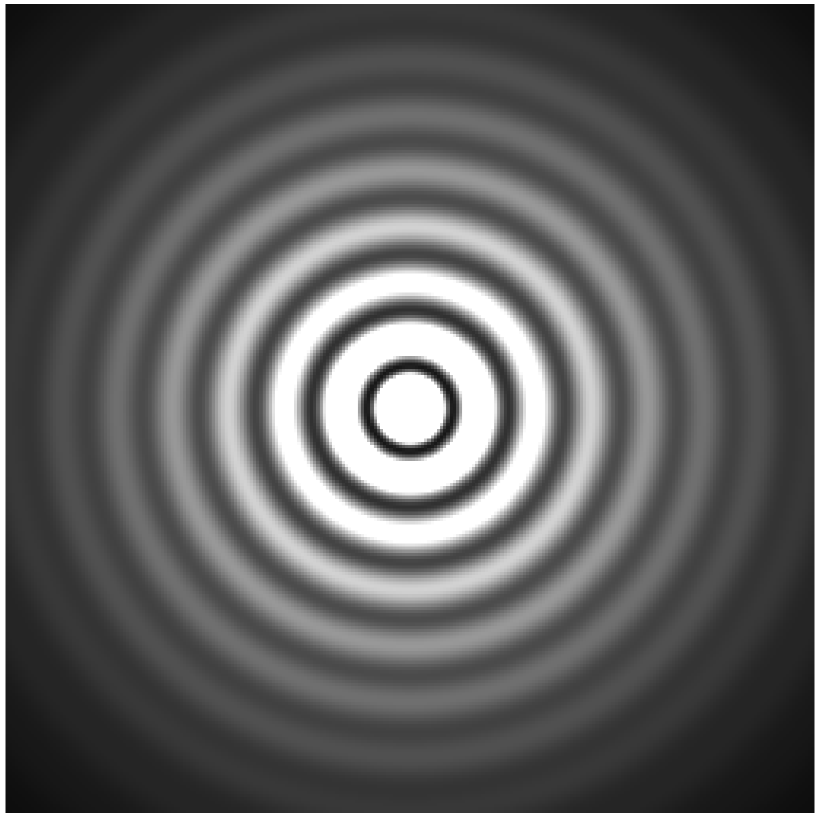

Moiré and Bessel

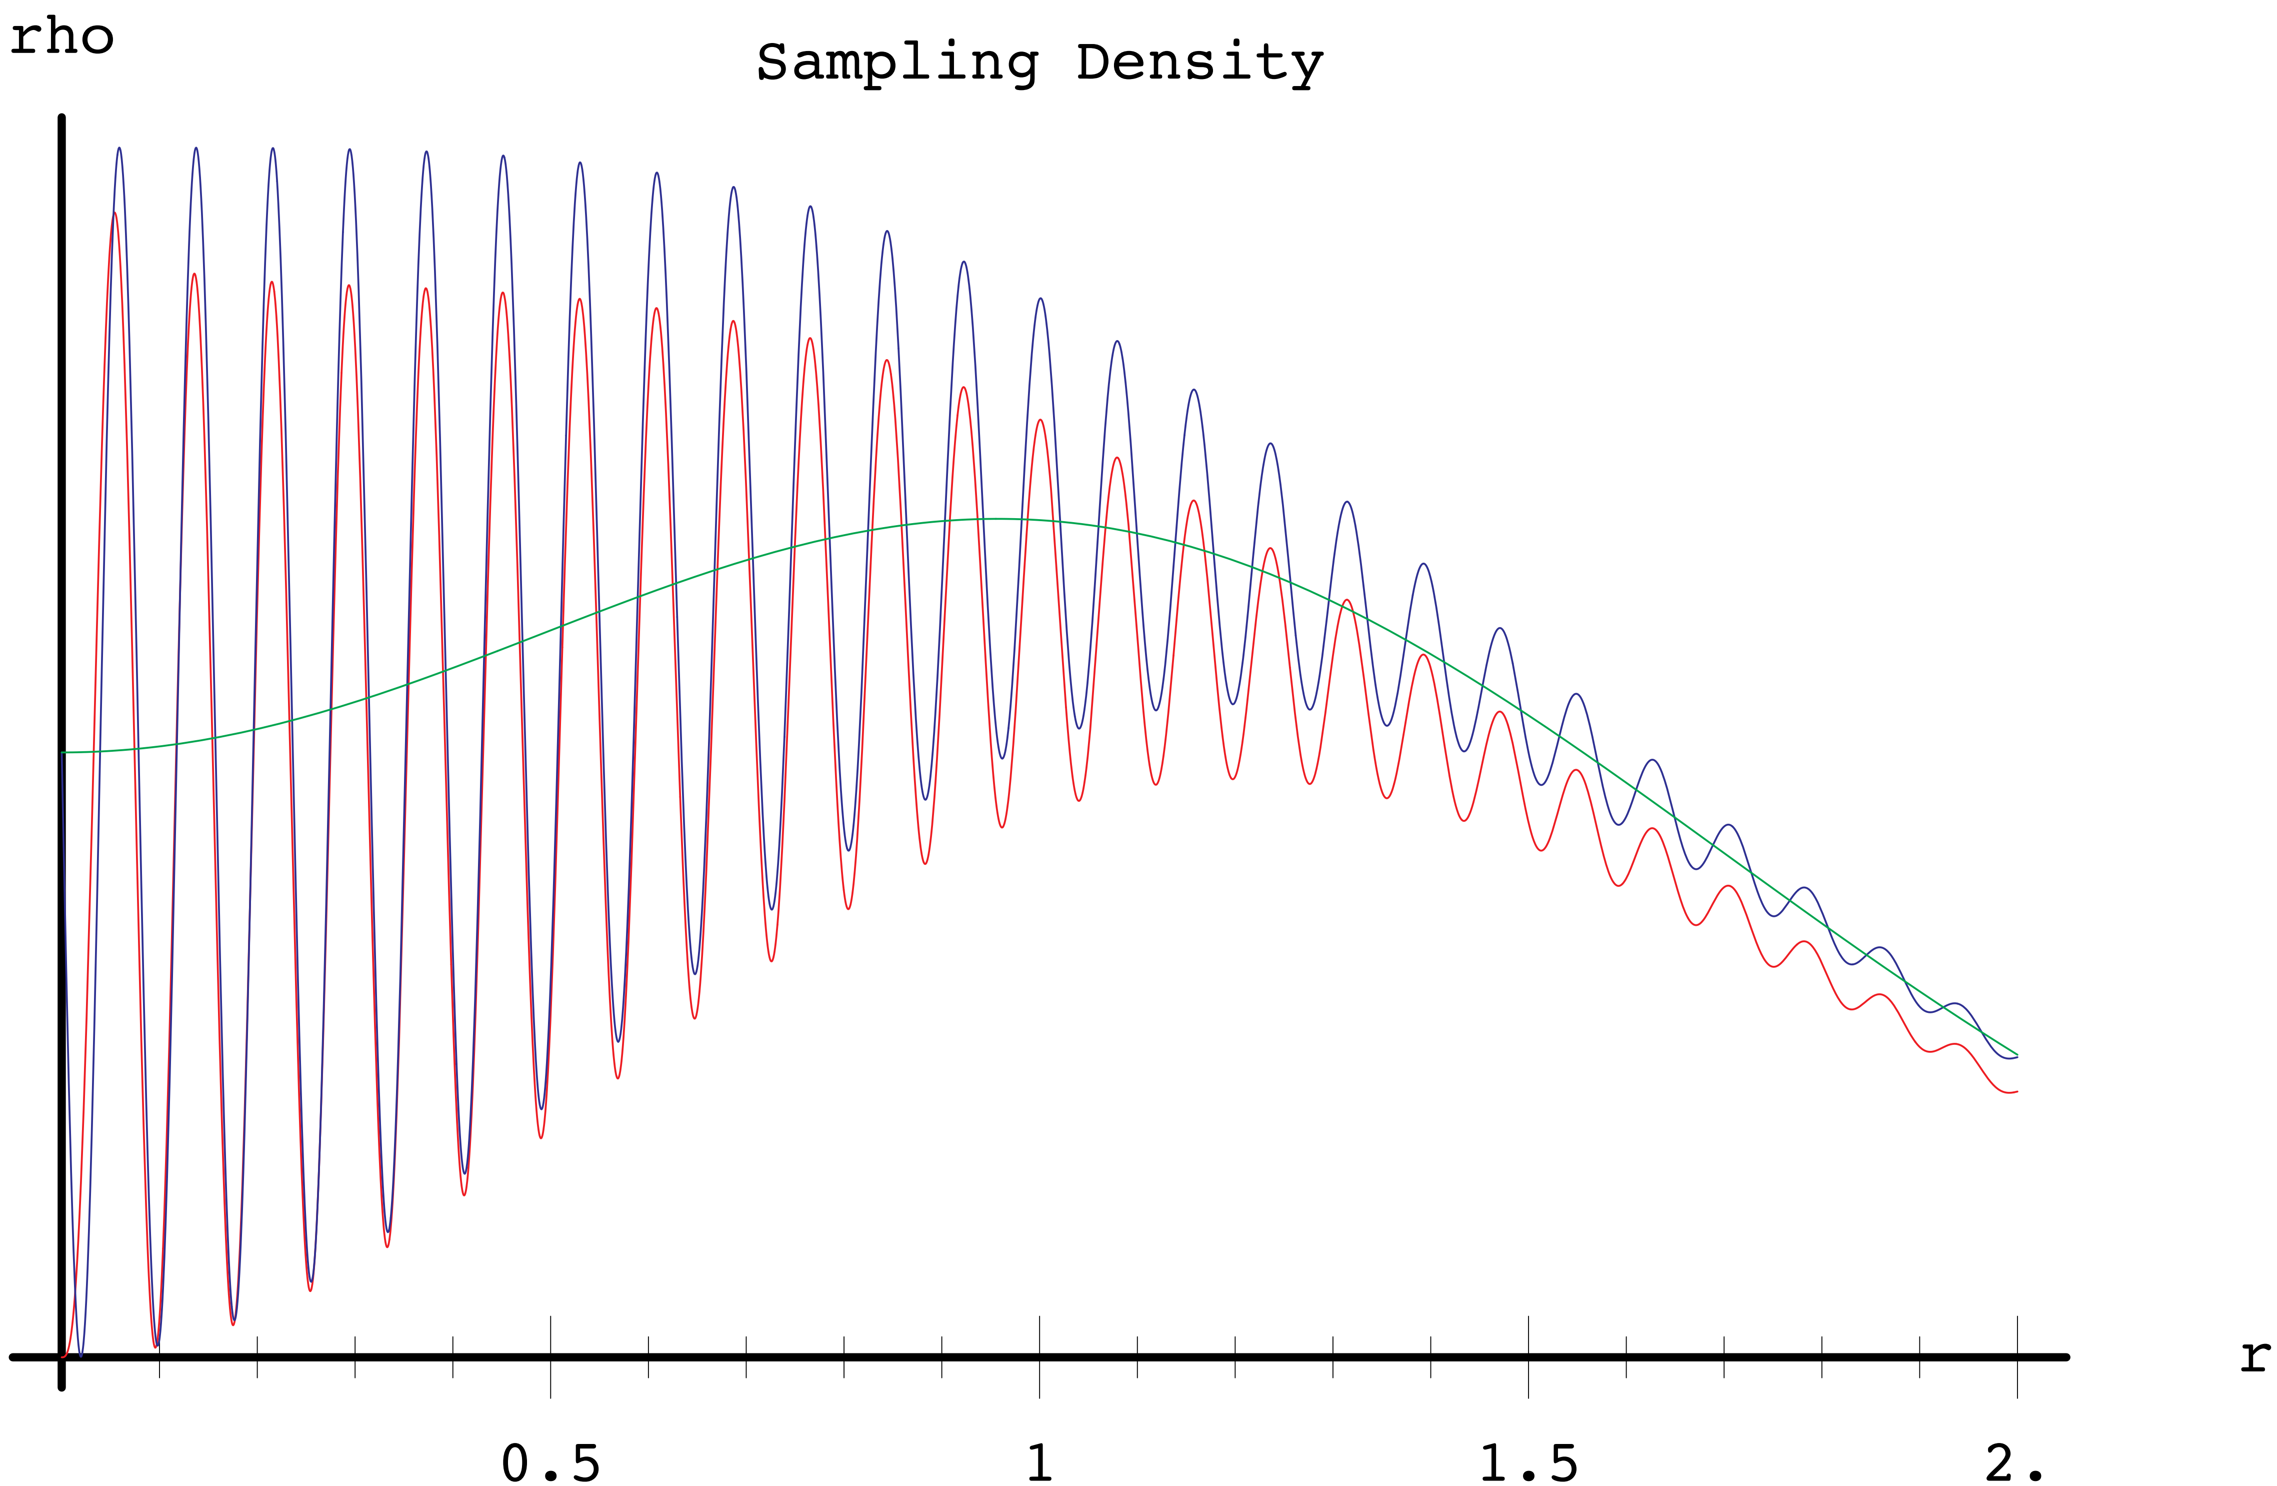

Combining multiple samples

Each sample has $r < \varepsilon$ with probability

$$\Pr[r < \varepsilon] \sim \varepsilon$$

Toy model: each point is uniformly random on a sphere of radius *r*, where *r* is uniformly random in [0,1]

How can we use *s* independent samples to obtain an estimate of the center with a very small ε? And how many samples does this take?

The mean (or the median) gives

$$s \sim \varepsilon^{-2} \log n$$

which isn't any better than taking *s* samples from the original shell!

Maximum likelihood

Our estimate will be one of the samples

With *s* samples, the closest one probably is at distance

$$\varepsilon \sim 1/s$$

but how do we know which one is the closest?

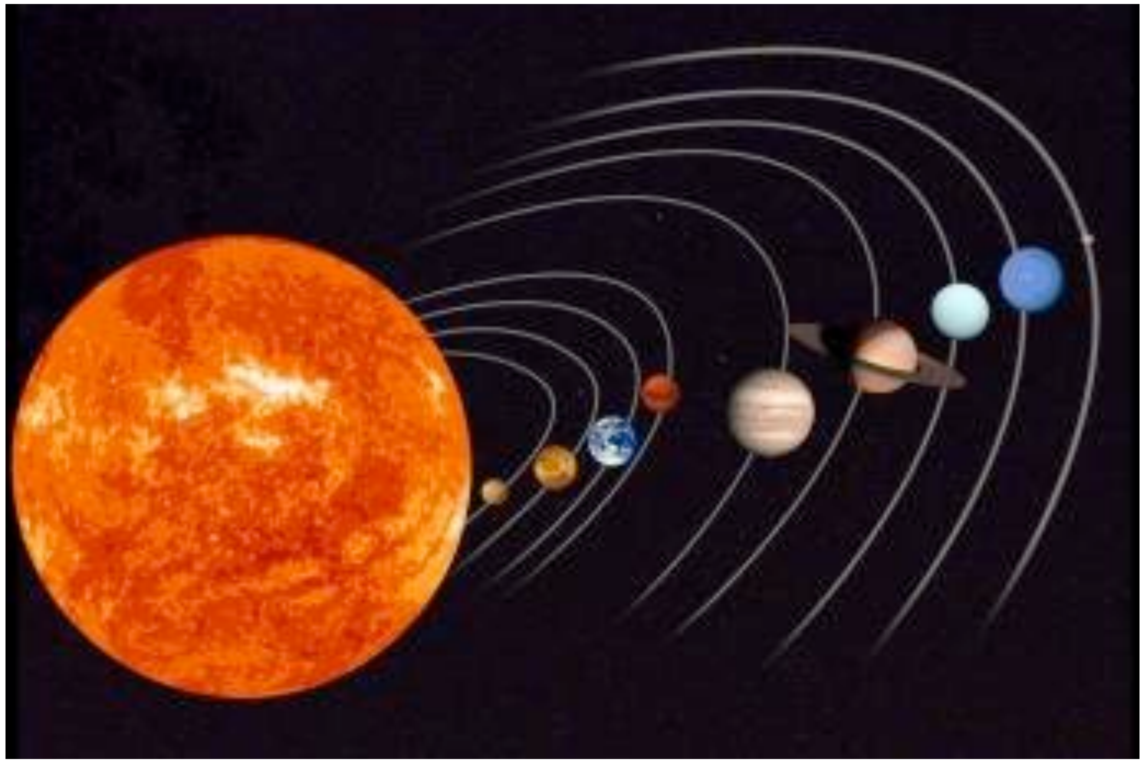

Maximum likelihood

Since $|\psi|^2 \sim r^{-(n-1)}$, if x_i were the center, the probability that we see the other samples is proportional to

$$L = \prod_{j \neq i} \left| x_i - x_j \right|^{-(n-1)}$$

We renormalize away the “self-likelihood” of x_i

Maximizing L is equivalent to minimizing the product of the distances

$$\prod_{j \neq i} \left| x_i - x_j \right|$$

Theorem: with probability bounded above zero, the x_i minimizing this product is the closest one

Gives error ε with just $s \sim 1/\varepsilon$ samples

Iterating to estimate the center of a spherically symmetric function

In one round of s samples, we get an estimate with error $\varepsilon \sim 1/s$

Create a uniform superposition in a ball around that estimate, sample in that ball to get a new estimate

But this ball has to have radius $n^{1/2}\varepsilon$ so the sampled shell has a large solid angle

Optimum scheme uses $en^{1/2}$ samples in each round, so that the ball's radius decreases by a factor of e at each step

Total number of samples over all log ε^{-1} rounds is then

$$s \sim n^{1/2} \log \varepsilon^{-1}$$

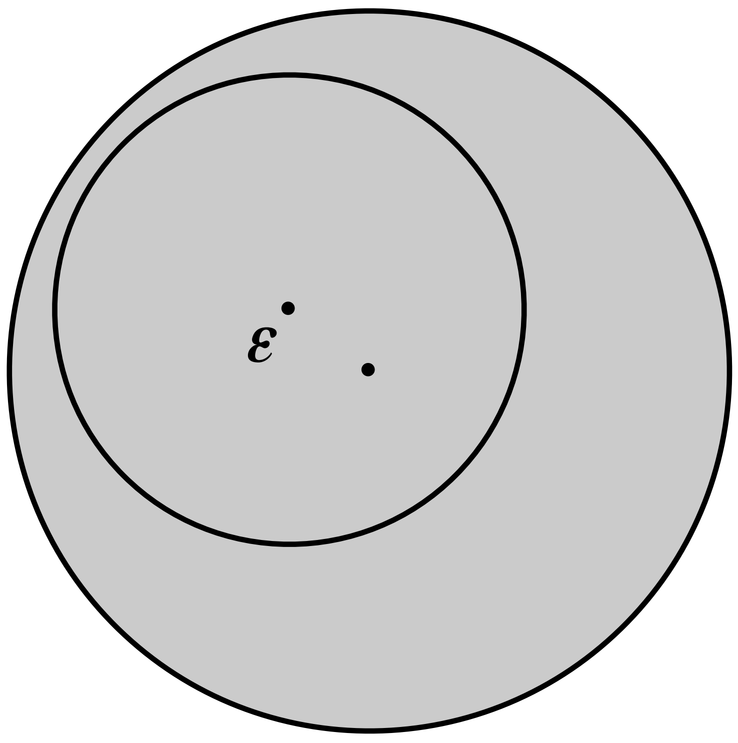

which is better than $s \sim 1/\varepsilon$ if $\varepsilon \ll n^{-1/2}$

Yi-Kai Liu's algorithm

Apply the *curvelet transform* to the sphere

Gives an approximate vector normal to a random point on the surface

One measurement: choose a random point along that vector

Two measurements: find intersection between two such vectors

Estimates the center within error

$$\varepsilon \sim (n w_0)^{1/2}$$

with just two queries

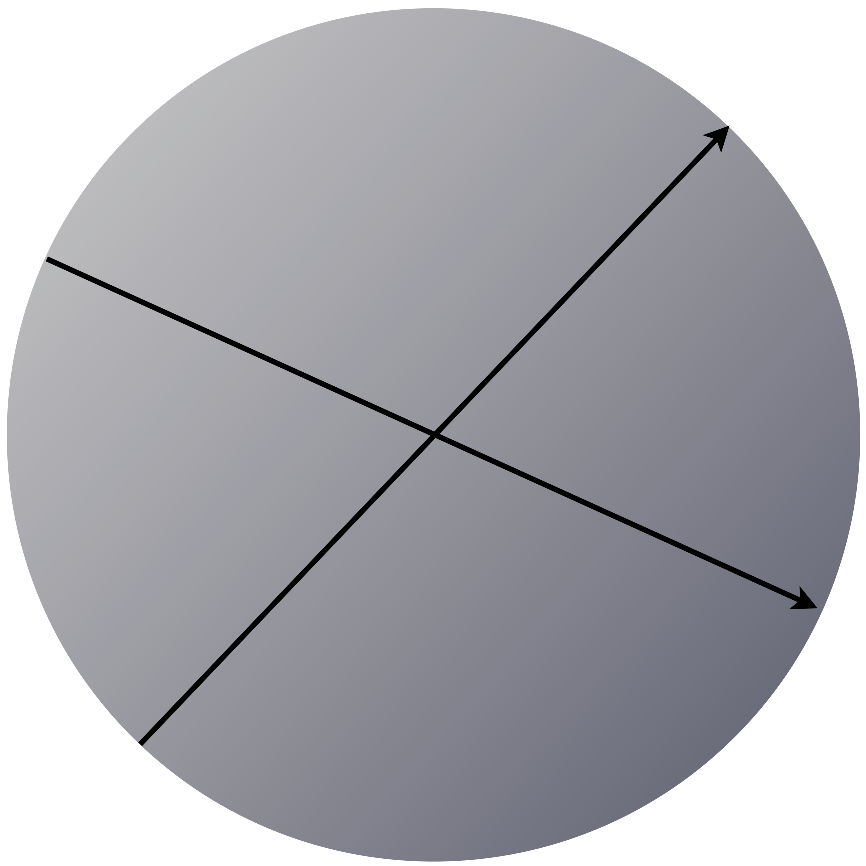

Other initial shapes? The evolute

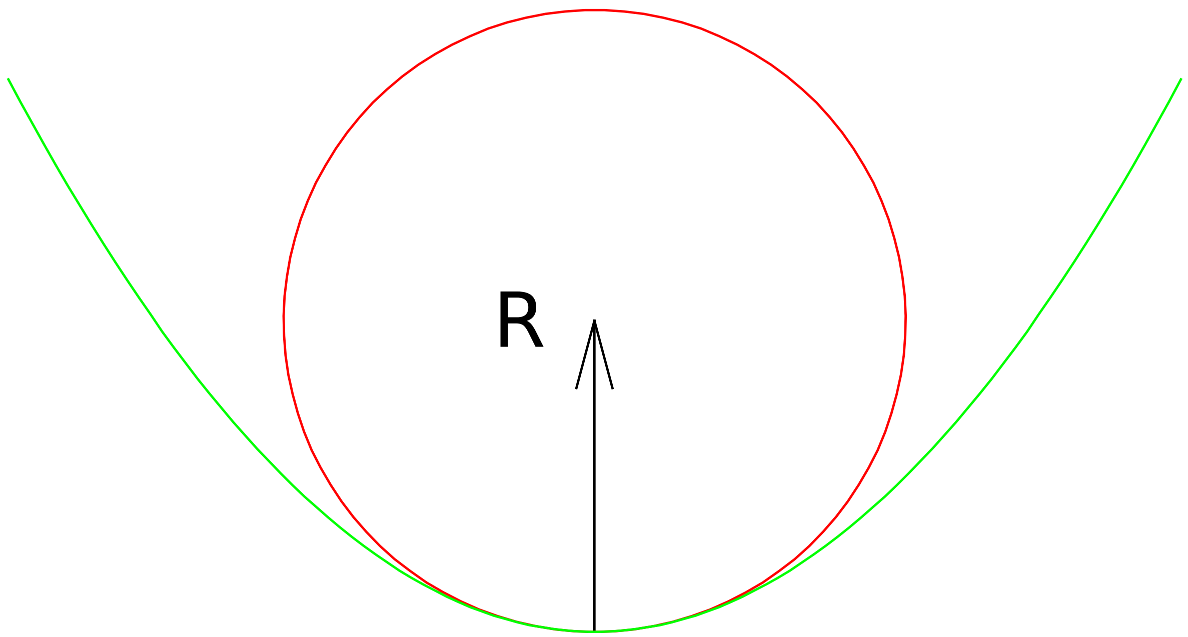

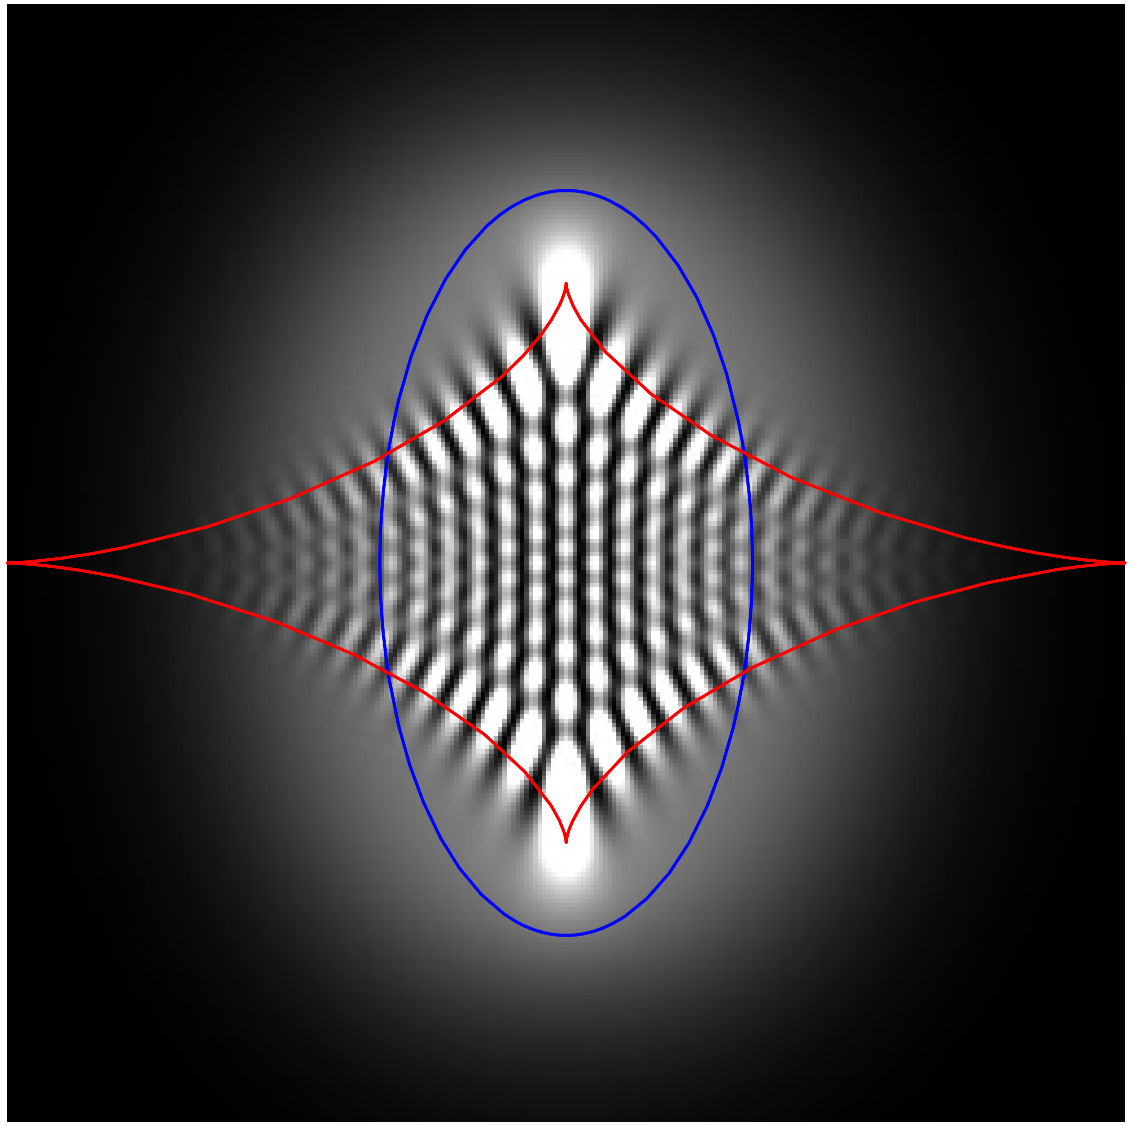

Scaling of the probability density

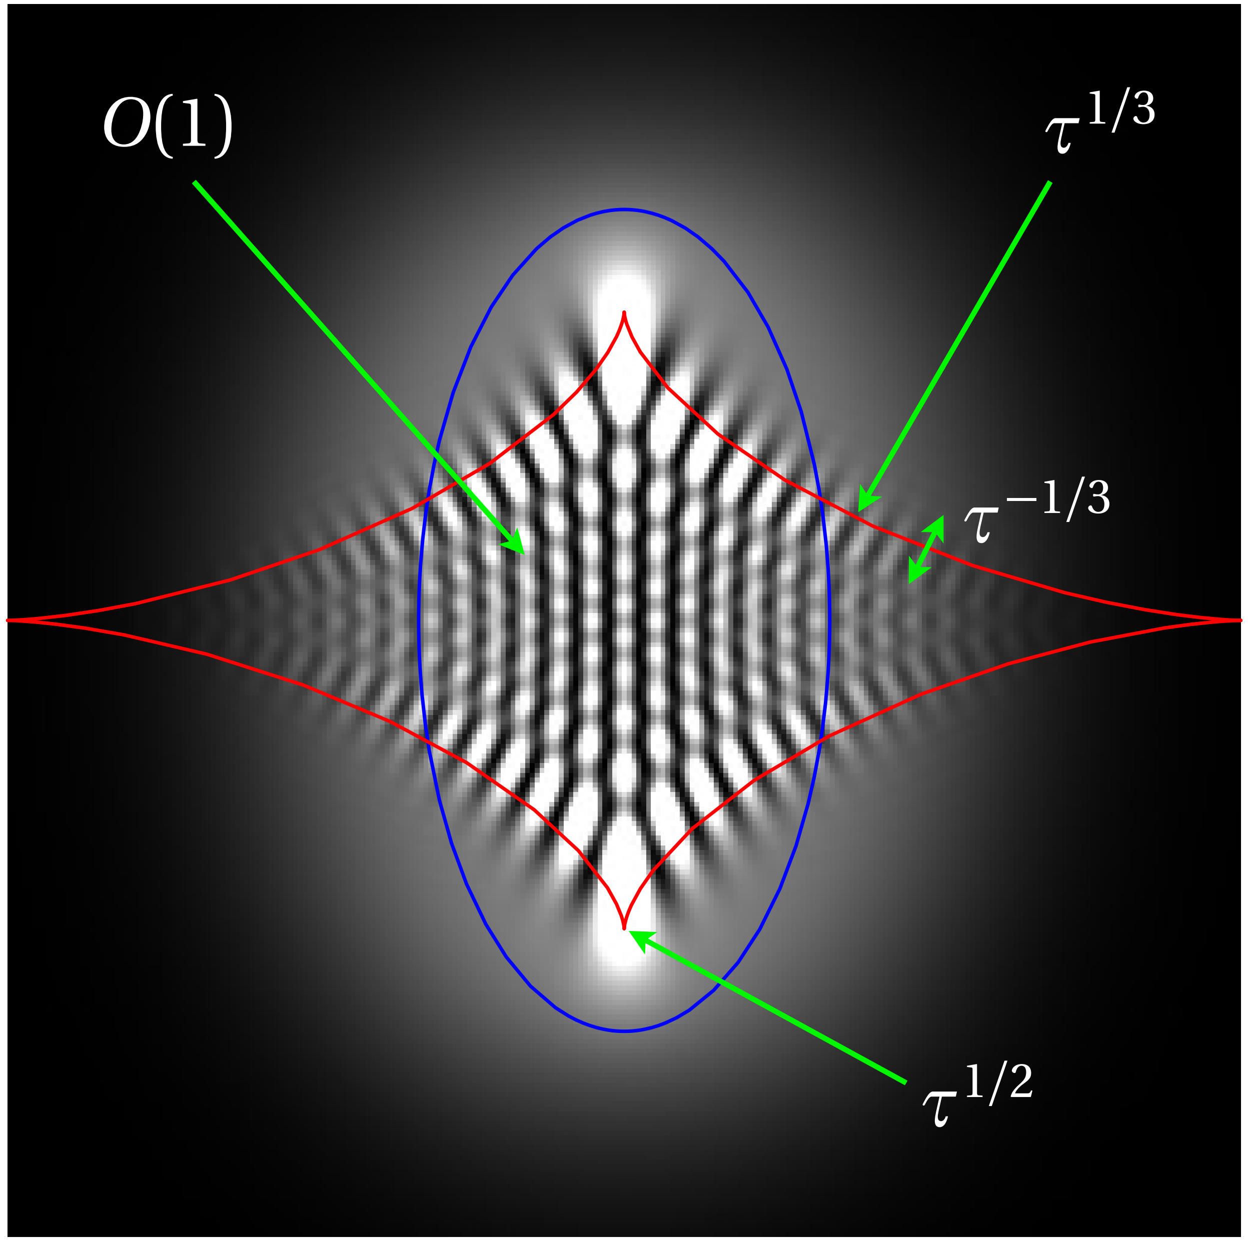

Computation everywhere

Even very simple physical processes can solve interesting problems, if they take place in unusual settings

The input to an algorithm can be given in the form of a Hamiltonian, a unitary operator, or an initial state

What other serendipities are still out there?

Shameless Plug

This book rocks! You somehow manage to combine the fun of a popular book with the intellectual heft of a textbook.

— Scott Aaronson

A treasure trove of information on algorithms and complexity, presented in the most delightful way.

— Vijay Vazirani

A creative, insightful, and accessible introduction to the theory of computing, written with a keen eye toward the frontiers of the field and a vivid enthusiasm for the subject matter.

— Jon Kleinberg

Oxford University Press, 2011

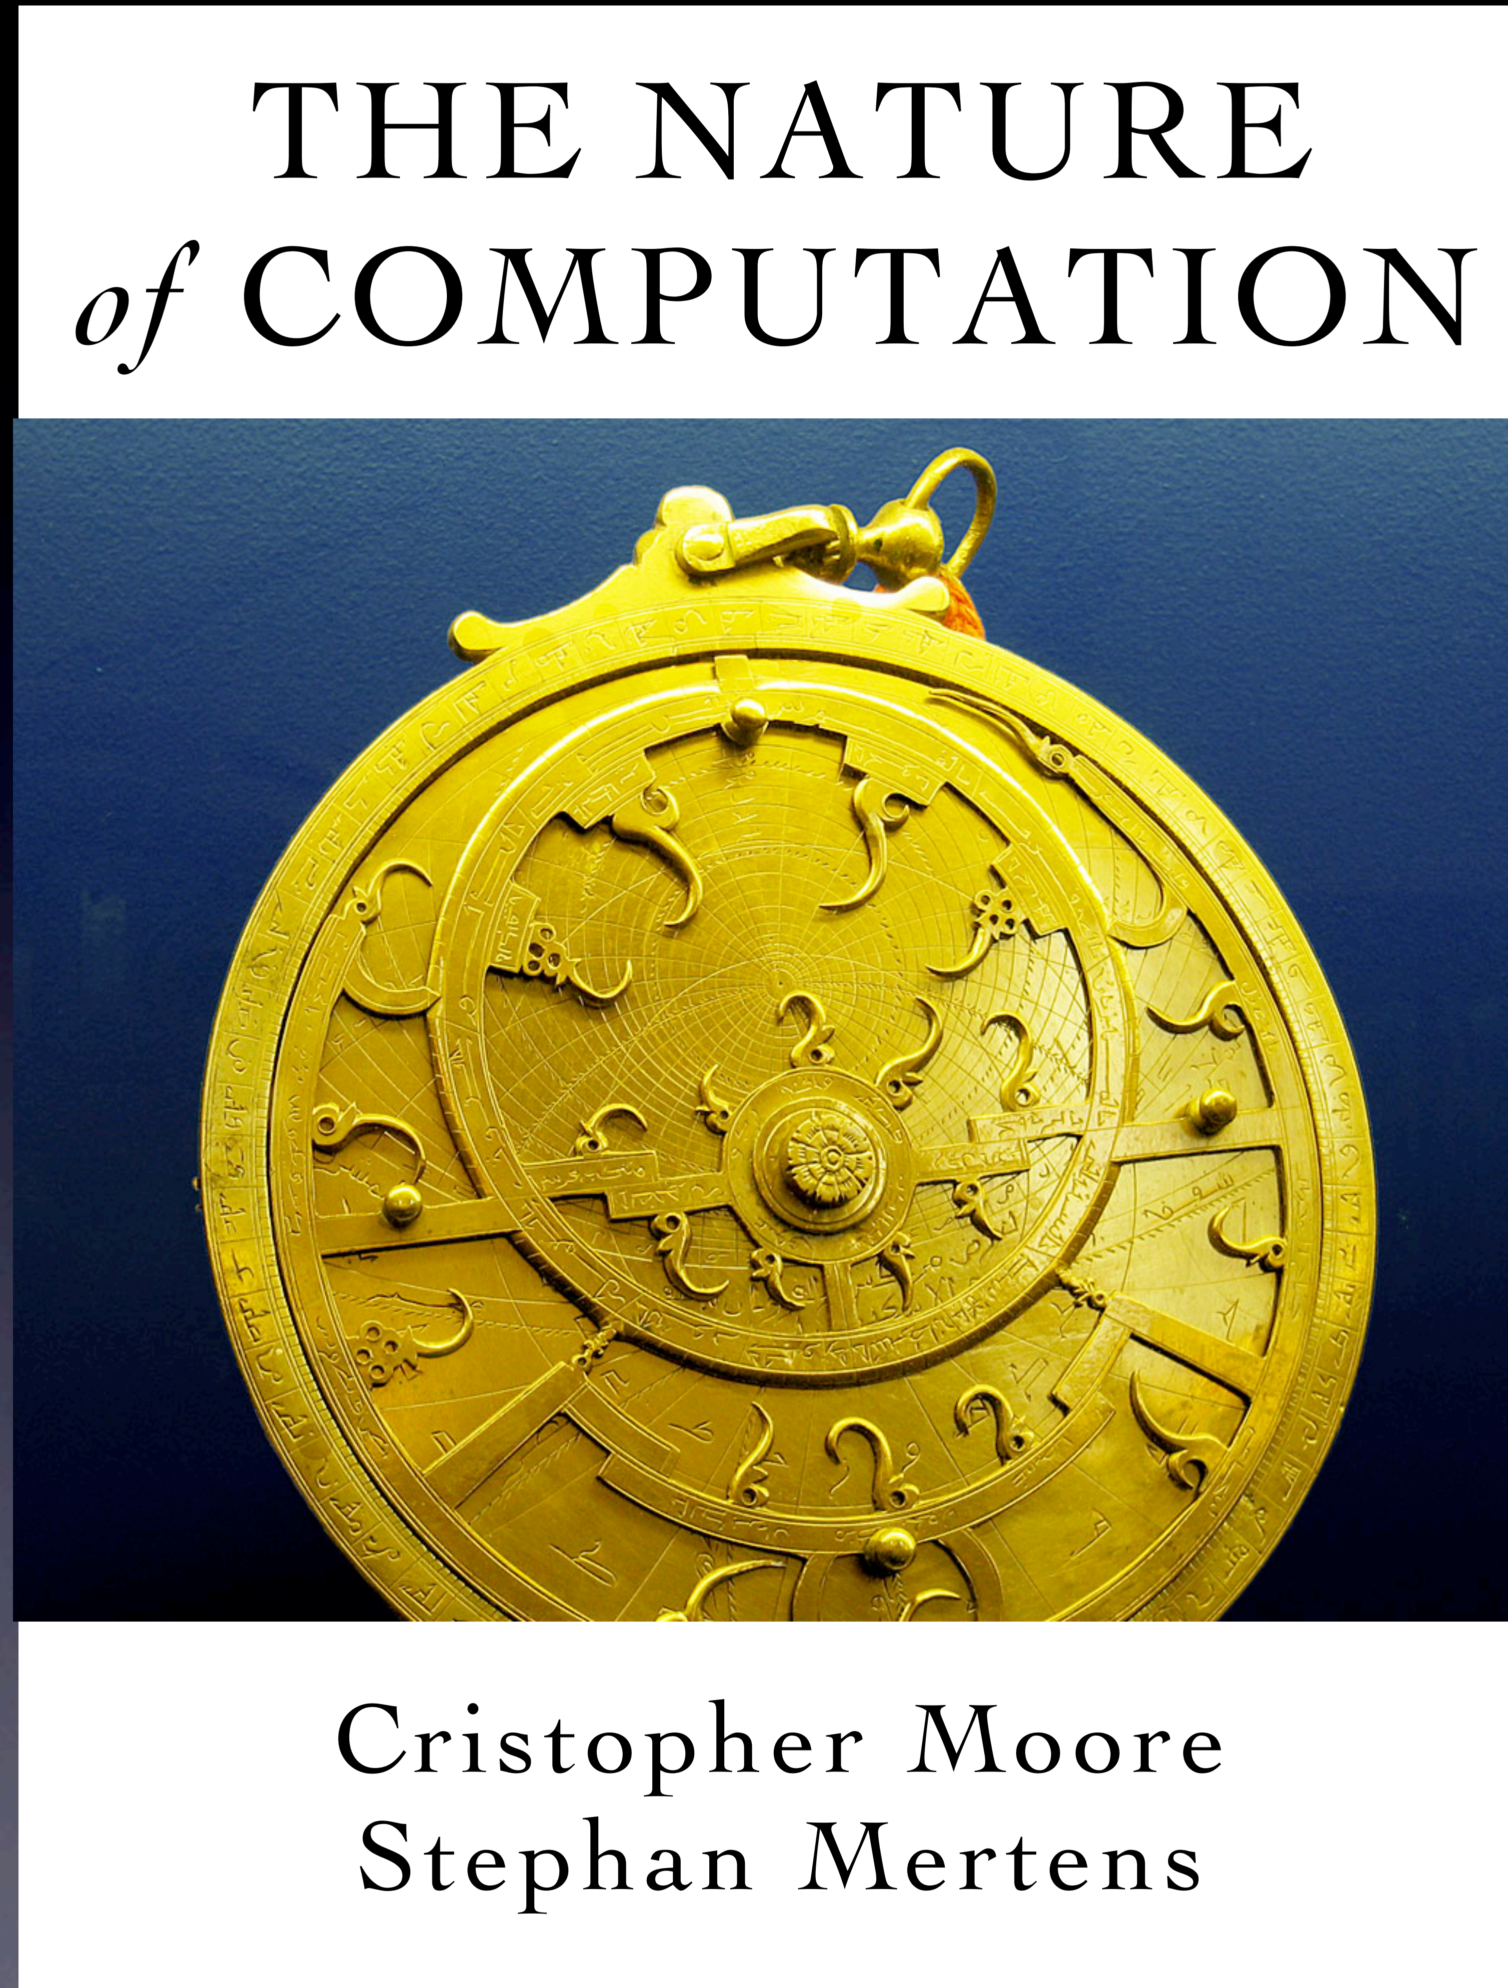

Acknowledgments

and the NSF and NSA/ARO