Deep Learning over Code.
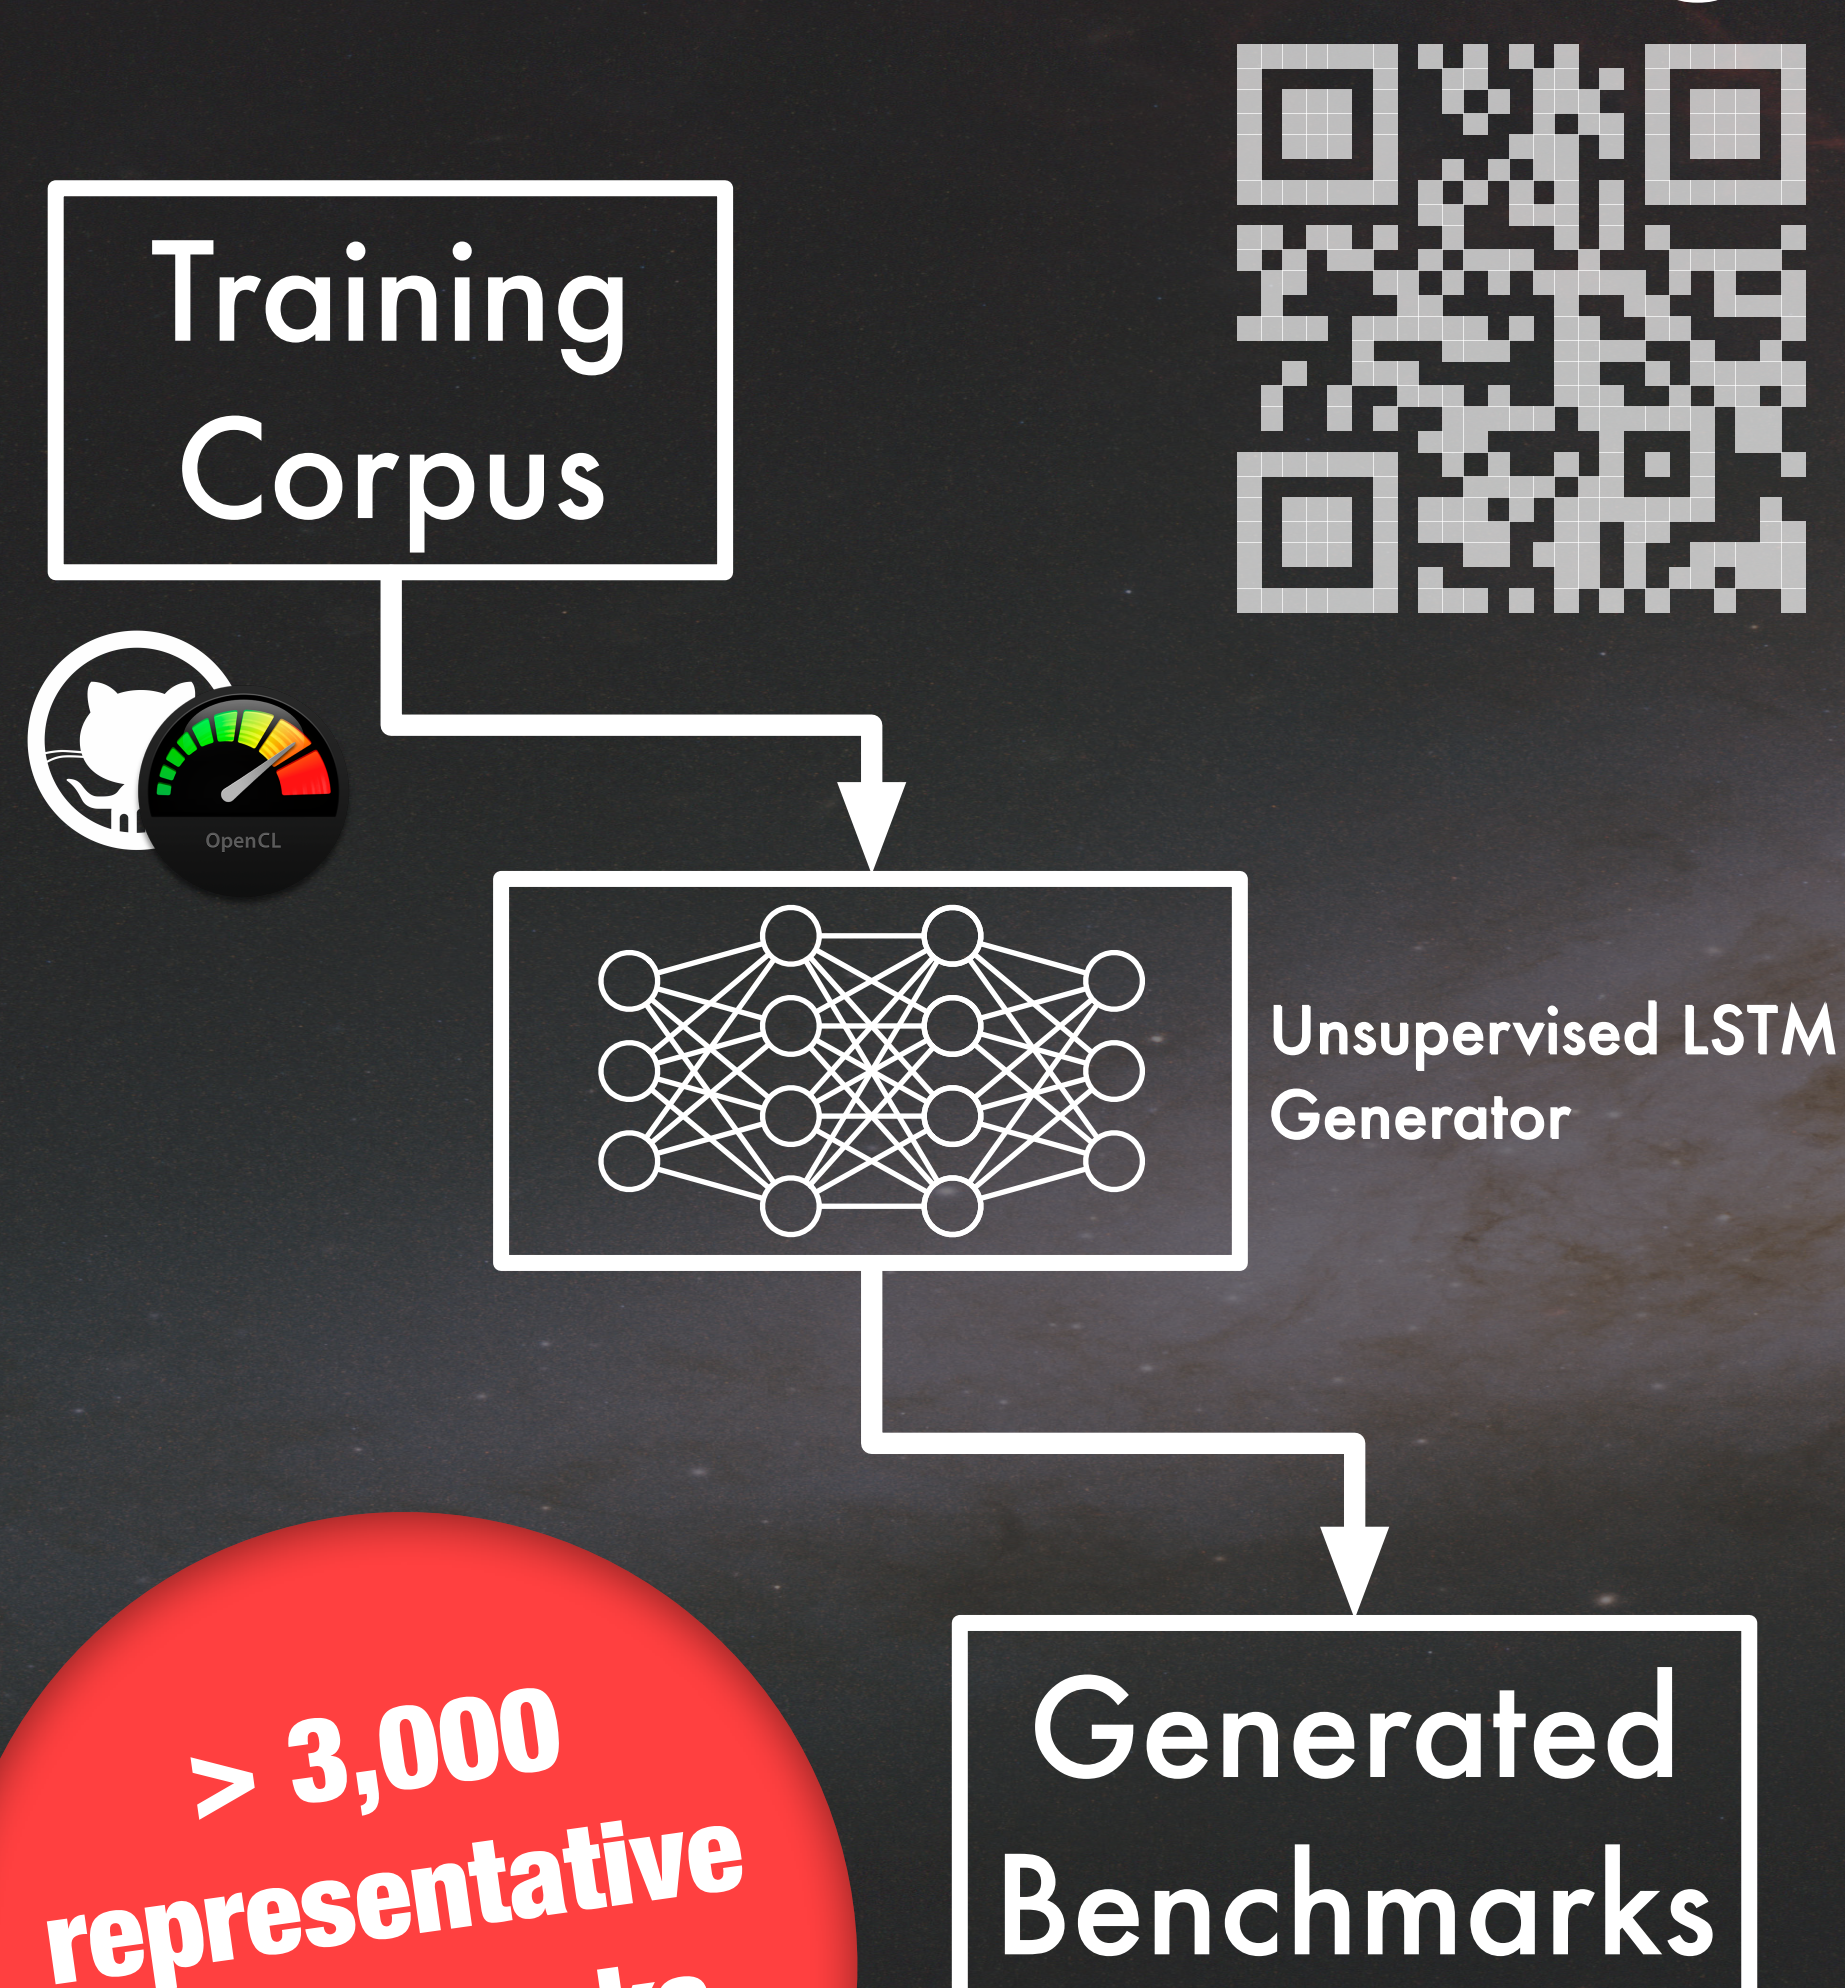
Chris Cummins, Pavlos Petoumenos, Alastair Murray, Zheng Wang, Hugh Leather.

c.cummins@ed.ac.uk

Best Paper CGO'17
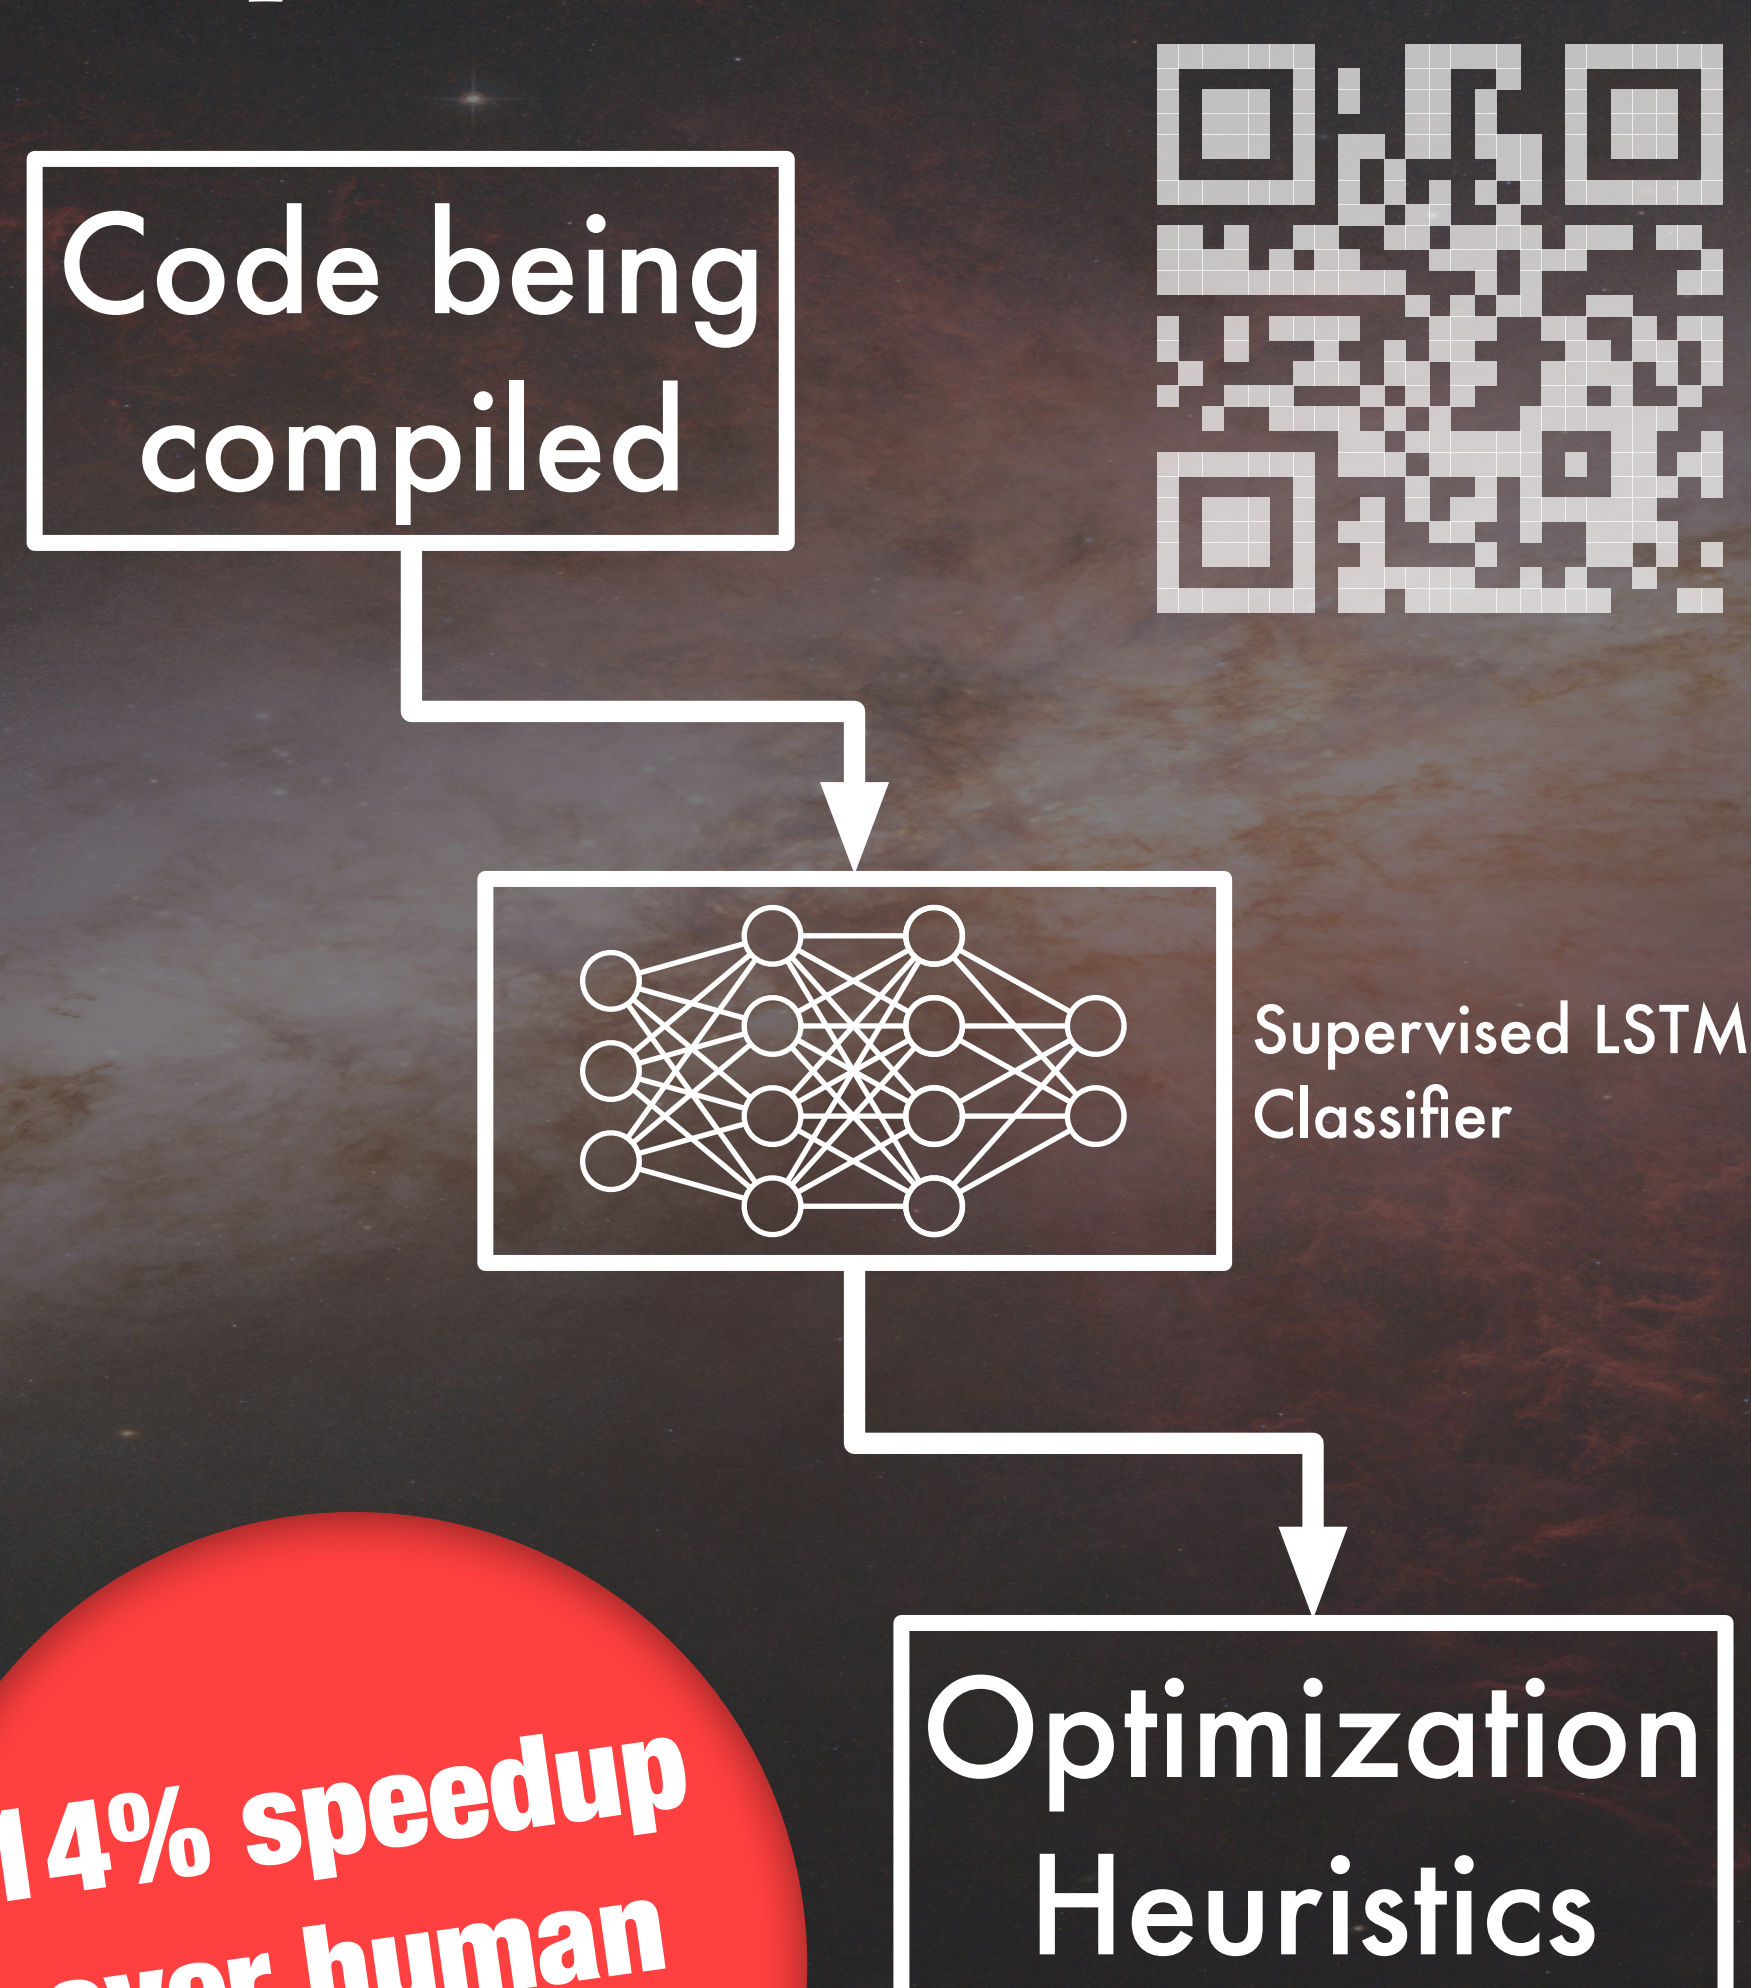
Best Paper PACT'17

Distinguished Paper ISSTA'18

Synthesizing Benchmarks for Predictive Modeling

Training Corpus → Unsupervised LSTM Generator → Generated Benchmarks

> 3,000 representative benchmarks per day

End-to-end Deep Learning of Optimization Heuristics

Code being compiled → Supervised LSTM Classifier → Optimization Heuristics

14% speedup over human expert

Compiler Fuzzing through Deep Learning

Training Corpus → Unsupervised LSTM Generator → Compiler Tests
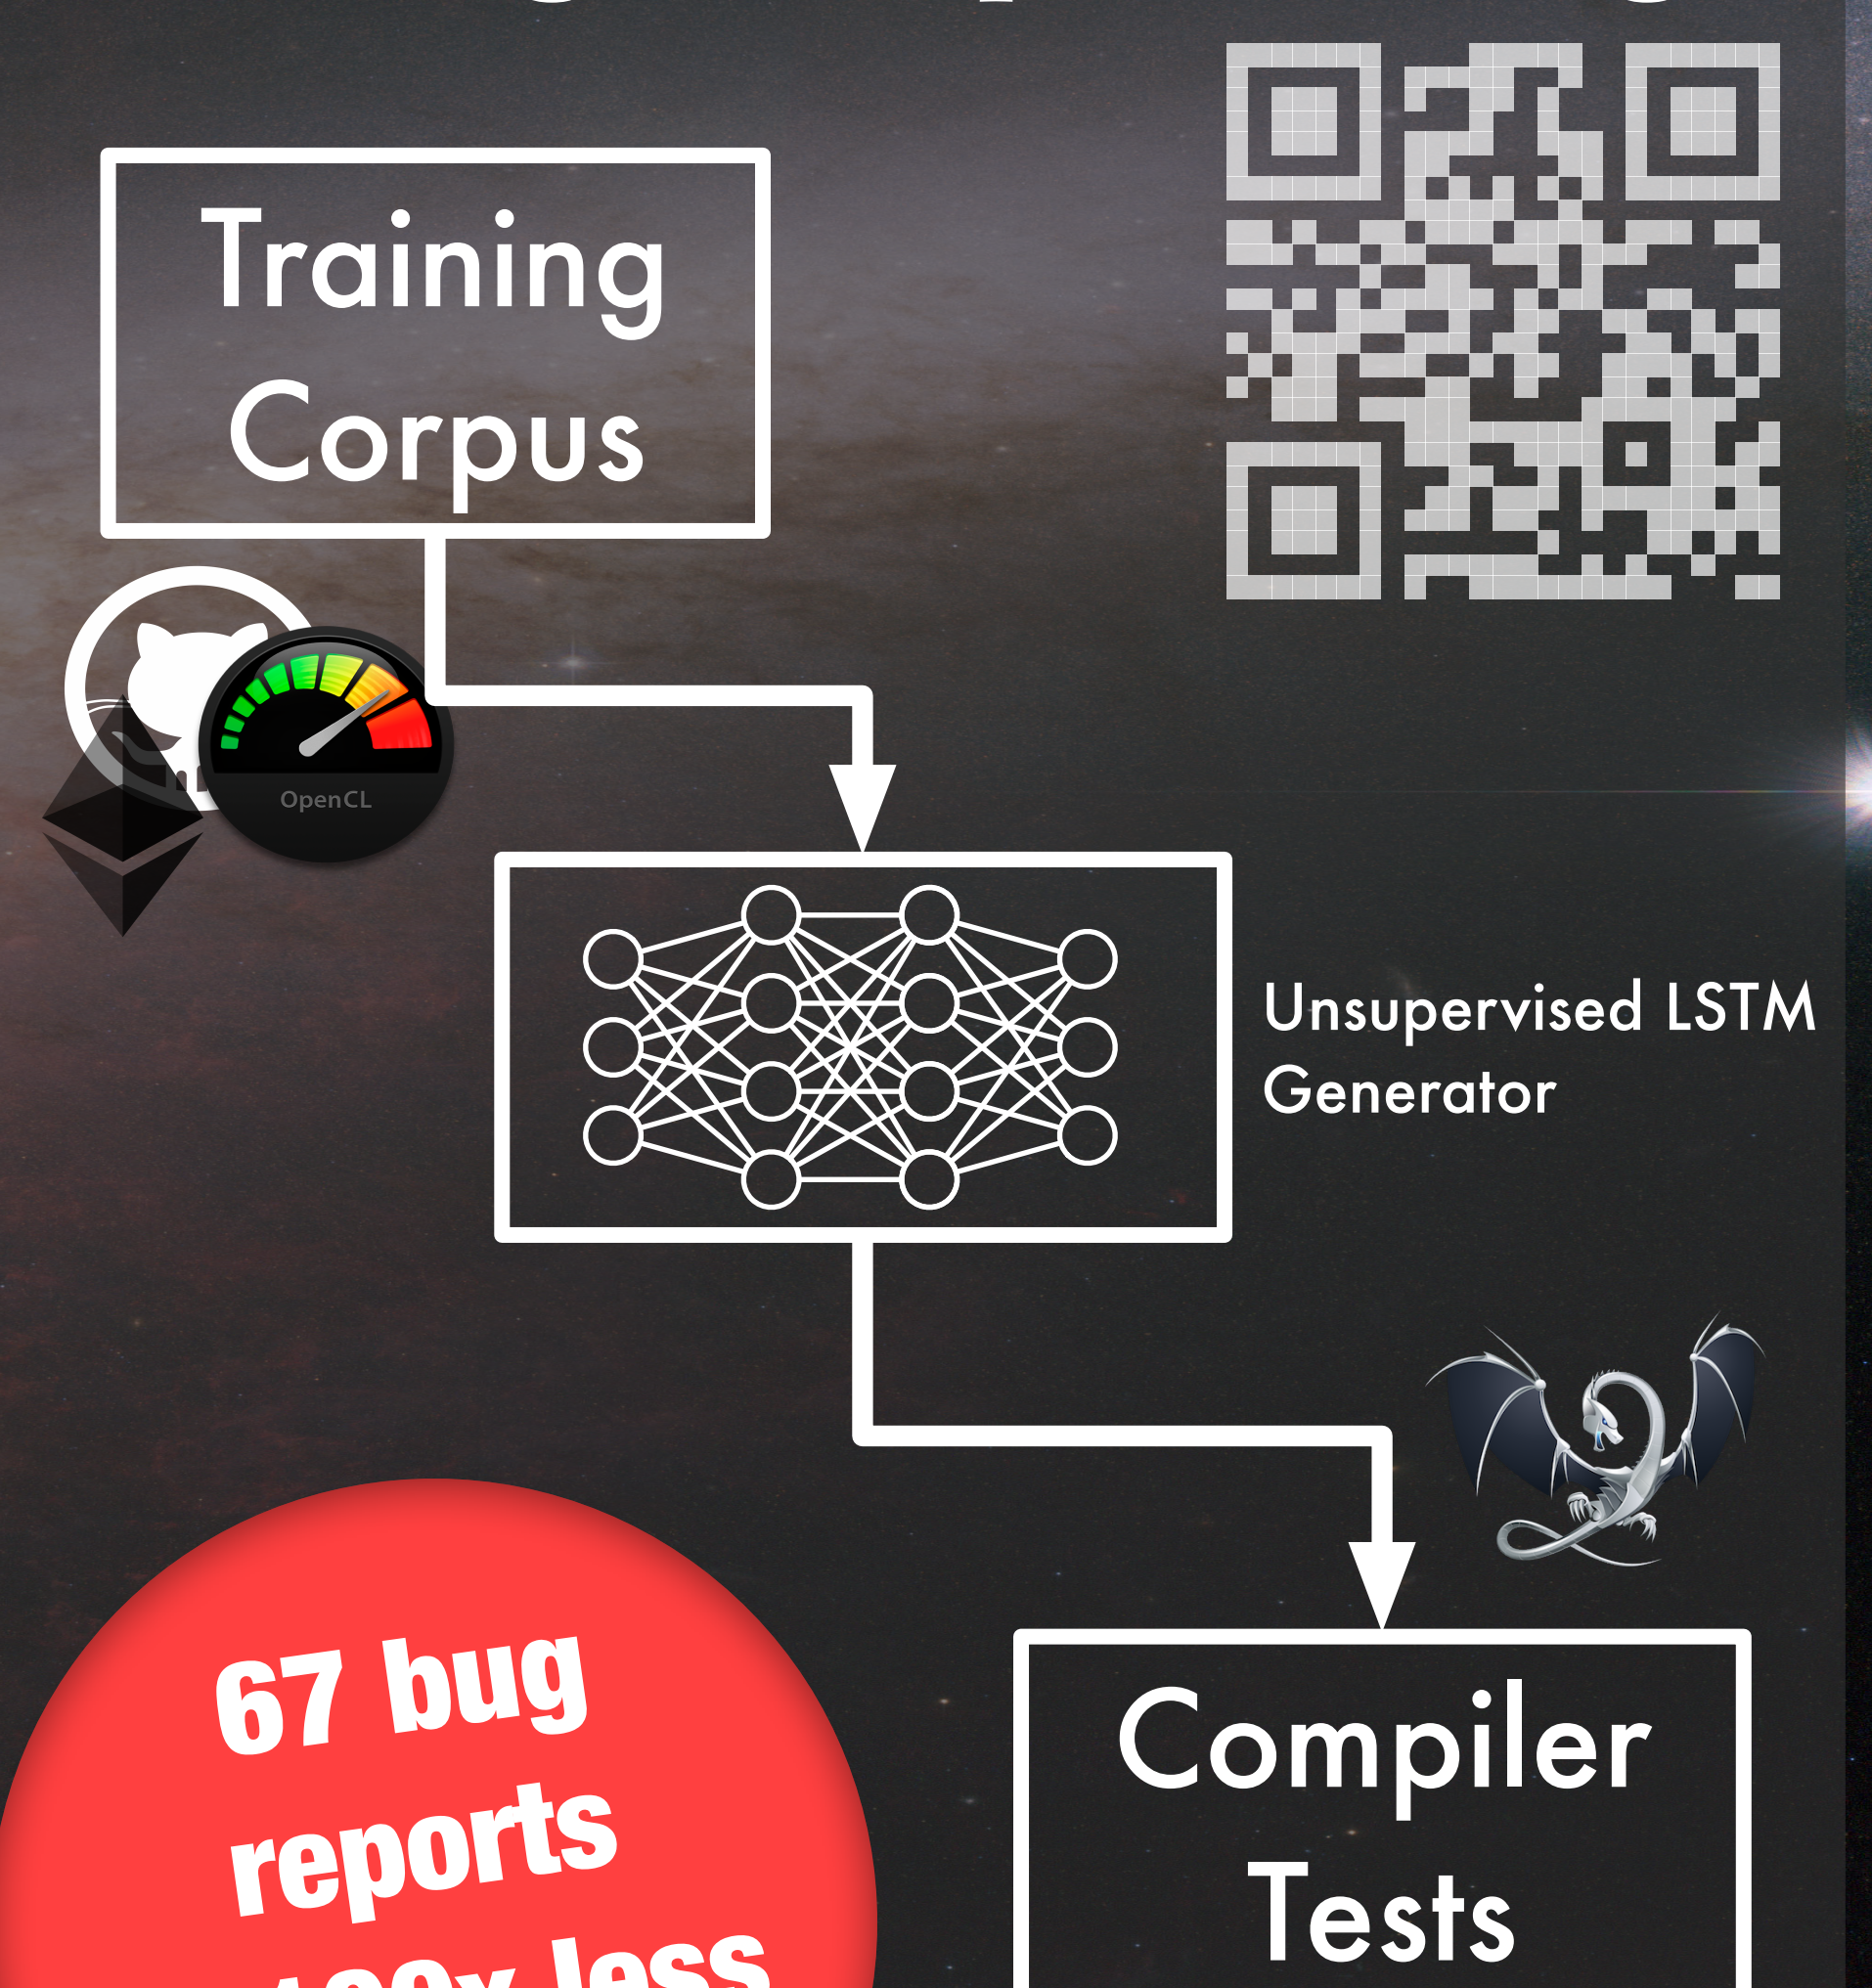
67 bug reports in 100x less code