


Nine Ways to Bias Open-Source AGI Toward Friendliness

Ben Goertzel

CEO, Novamente LLC and Biomind LLC
CTO, Genescent Corp
Co-founder, OpenCog Project
Vice Chairman, Humanity+
Adjunct Research Professor, Xiamen University, China
Advisor, Singularity University and Singularity Institute

Joel Pitt

M-Lab, Hong Kong Polytechnic University
Co-founder, OpenCog Project



1

Engineer the capability to acquire
integrated ethical knowledge

2

Provide rich ethical interaction
and instruction, respecting
developmental stages

3

Develop stable, hierarchical goal
systems

4

Ensure that the early stages of recursive self-improvement occur relatively slowly and with rich human involvement

5

Tightly link AGI with the Global
Brain

6

Foster deep, consensus-building
interactions between divergent
viewpoints

7

Create a mutually supportive
community of AGIs

8

Encourage measured co-advancement
of AGI software and AGI ethics
theory

9

Develop advanced AGI sooner not
later

A paper corresponding to this presentation,
reviewing each of the 9 points in detail,
is in preparation
and will be submitted for publication this fall