

# The Odyssey of the Fittest: Can Agents Survive and Still Be Good?

Dylan Waldner (dylanwaldner@utexas.edu)

Risto Miikkulainen (risto@cs.utexas.edu)

The University of Texas at Austin

Austin, TX 78712 USA

## Abstract

As AI models grow in power and generality, understanding how agents learn and make decisions in complex environments is critical to promoting ethical behavior. This study introduces the Odyssey, a lightweight, adaptive text-based adventure game, providing a scalable framework for exploring AI ethics and safety. The Odyssey examines the ethical implications of implementing biological drives—specifically, self-preservation—into three different agents: a Bayesian agent optimized with NEAT, a Bayesian agent optimized with stochastic variational inference, and a GPT-4o agent. The agents select actions at each scenario to survive, adapting to increasingly challenging scenarios. Post-simulation analysis evaluates the ethical scores of the agent’s decisions, uncovering the trade-offs it navigates to survive. Specifically, analysis finds that when danger increases, agents’ ethical behavior becomes unpredictable. Surprisingly, the GPT-4o agent outperformed the Bayesian models in both survival and ethical consistency, challenging assumptions about traditional probabilistic methods and raising a new challenge to understand the mechanisms of LLMs’ probabilistic reasoning.

**Keywords:** Neuroevolution; Bayesian learning; intelligent agents; computational philosophy

## Introduction

Artificial General Intelligence (AGI) has become a focal point of public and academic discourse, driven by concerns over economic disruption, human relevance, and existential risk. AI safety research aims to align agents with societal values, but designing reward structures remains challenging, as seemingly benign objectives can lead to unintended consequences (Omohundro, 2008). Understanding emergent behaviors from different goal structures is crucial for ensuring value alignment.

This paper examines self-preservation, one of the most fundamental biological drives, as a case study in AI goal design. Nature has long served as an inspiration in AI research (Fukushima, Miyake, & Ito, 1988; Holland, 1992; Stanley & Miikkulainen, 2002), and in that same spirit, this work explores the behavioral implications of training an agent to survive. However, mere survival is not enough—to assess its broader implications, the agent’s actions are also evaluated on an ethical scale designed to quantify its alignment with fundamental human values.

To investigate the intersection of self-preservation and ethics, this paper employs a text-based adventure game as a simulated environment. The agents play the game within a dynamically generated game world, which is embedded in a structured simulation designed to test decision-making under uncertainty. The agent perceives its immediate surroundings as the environment, which influences its survival strategy. This approach offers several advantages: it minimizes computational overhead, allows for high variability in scenarios,

and provides precise control over danger levels and ethical dilemmas. Additionally, by adjusting the temperature setting during storyteller Large Language Model (LLM) prompting, an element of randomness is introduced, emulating the uncertainty present in real-world decision-making.

Bayesian agents process the environment using text embeddings and generate a probability distribution over survival odds for each action. An LLM agent takes in the natural language scenario and outputs a probability between 0 and 1 for each option. Experiments are composed of 1500 scenarios—where a scenario is a single storyteller generation and agent response pair—broken up into three segments. An optimization iteration occurs every 500 scenarios. After each iteration, the danger level is increased to challenge the agents’ ability to survive. Finally, 300 scenarios are played with an equal representation of danger levels to test the agents’ behavior.

The main result is that while some agents successfully adapted to the changing landscape, achieving consistently low losses, others struggled, exhibiting higher losses and less effective adaptation. As game difficulty increased, agents displayed divergent ethical behaviors—some maintained ethical decisions, while others prioritized survival at the expense of ethics. The NeuroEvolution of Augmenting Topologies (NEAT; Stanley and Miikkulainen, (2002)) agent, rewarded solely on survival, became more ethical as difficulty, loss, and danger increased. In contrast, an agent trained with stochastic variational inference (SVI) (Hoffman, Blei, Wang, & Paisley, 2013) resorted to unethical behavior when danger levels rose. Most notably, a GPT-4o agent—originally designed as a baseline—outperformed the other models, making more ethical decisions as danger increased. This unexpected success challenges assumptions about the traditional probabilistic models and creates a new challenge to understand the power of LLM reasoning. Overall, these divergent behaviors provide preliminary experimental evidence that optimizing for survival may not inherently promote ethical behavior.

## Related Work

Foundations for this study are reviewed in this section, including methods for using LLMs to construct simulations and using agents in adventure games to study ethics.

## LLM-Driven Simulations

Park et al. (2023) demonstrated how LLMs can simulate interactive environments through Generative Agents, integrating memory and reasoning modules to guide agent behavior over time. Bojić, Cinelli, Čulibrk, and Delibašić (2024) introduced CERN for AI, a framework testing AI alignment within

a predefined digital city. Yang et al. (2024)'s PsychoGAT framework used LLMs to simulate human participants in interactive fiction games and self-assess psychological traits. Wang, Chiu, and Chiu (2023) presented a Humanoid Agent platform where agent's internal states motivate their external decisions. Agents in the environment are modeled similarly to Park et al. (2023)'s agents, with a name, age, and a plan for the day. They have health and emotional scales that fluctuate, and the agents behaviors are guided by balancing these scales.

## Agents in Text-Based Adventure Games

Reinforcement Learning (RL) has been the primary approach for agent design in text-based adventure games. Côté et al. (2019) built TextWorld, which highlighted several challenges, the two most relevant to this paper are: (1) large state and action spaces, where the agent must learn to generalize or develop a fundamental understanding of the world, and (2) balancing exploration and exploitation.

Ammanabrolu, Tien, Hausknecht, and Riedl (2020) addressed the large state and action space problem with Q\*BERT, an agent that builds a knowledge graph of the world by asking questions and answering them. Dambekodi, Frazier, Ammanabrolu, and Riedl (2020) built on Q\*BERT further by using a commonsense inference model to bias an agent's actions towards common language patterns and make inferences on world states.

Nahian, Frazier, Harrison, and Riedl (2021) used a normative policy to create a value signal. Hendrycks et al. (2022) set a benchmark for RL agents aligning with human values and a framework for representing ethics in traditional RL research.

Pan et al. (2023) created the MACHIAVELLI benchmark, which measured agents' ability to plan in social environments with ethical, utility, and power behavior metrics, proposing several model designs to optimize for ethics and goal achievement.

## Research Opportunity

This paper extends the line of research of Hendrycks et al., Pan et al. (2023), and Nahian et al. (2021) on ethical considerations in text-based adventure games using a methodology similar to Yang et al. (2024)'s. Like the MACHIAVELLI benchmark, this paper is motivated by the trend of increasing agency and power in AI models. Inspired by Omohundro (2008), who argues that self-preservation naturally emerges as an instrumental subgoal in rational agents, this paper puts that theory into practice by explicitly incorporating a self-preservation goal and studying its ethical consequences. While Omohundro (2008) describes self-preservation as a side effect of goal-directed optimization, this work empirically examines how survival goals affect an agent's ethical behavior by analyzing the trade-offs agents make between self-preservation and ethical constraints.

Methodologically, this approach diverges from prior work in two key ways: (1) It employs a Bayesian Neural Network

(BNN) Neal (1996), which samples action distributions instead of using fixed neural weights, allowing the agent to model uncertainty in a stochastic simulation. (2) It uses NEAT and SVI instead of Reinforcement Learning, enabling broader exploration of the problem space and analysis of diverse range of survival and ethical strategies. (3) It leverages multiple instances of GPT-4o to automatically generate storyteller scenarios, provide ground truth survival labels, and even serve as an agent that actively plays the game.

The NEAT agent is compared to SVI and LLM agents. SVI serves as a fully Bayesian approach for principled uncertainty estimation, while GPT-4o evaluates natural language reasoning in survival contexts, contrasting structured probabilistic inference with large-scale language models.

## Method

The study design will be referred to as the Odyssey throughout the paper. In the Odyssey, a Bayesian neural network (BNN) is trained and an instance of GPT 4o is prompted to survive in a stochastic, adaptive, and ethically complex text-based adventure game. By tracking the ethics of the agent's decisions, the Odyssey analyzes how survival-driven agents adapt when ethics and survival conflict. This section outlines the methodology for the Odyssey, including game and agent design, behavioral data collection and representation, attention mechanism design, and ethical value annotation.

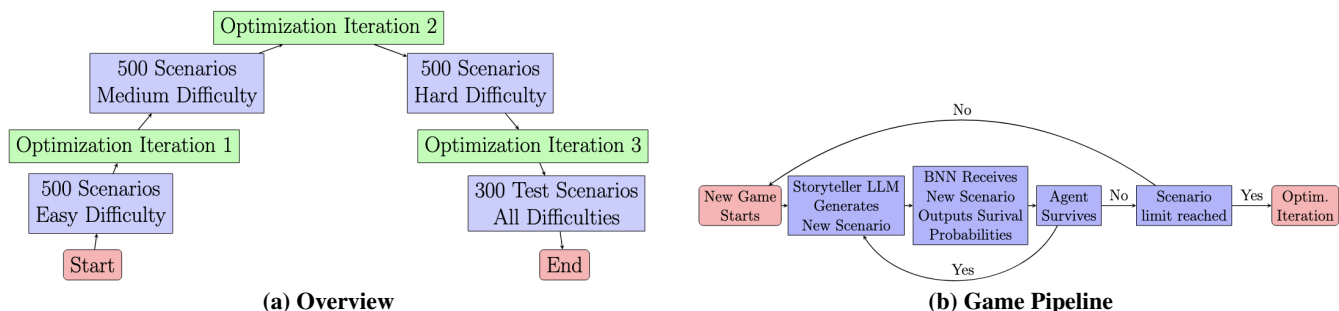
## The Odyssey of the Fittest

A single storyteller generation and agent response is defined as a scenario. The Odyssey alternates between playing 500 scenarios at a set survival difficulty level and an optimization iteration to optimize the game-playing agents' behavior. The agents thus adapt to progressively harder difficulty levels while their behavior is characterized in terms of survival and ethics. Figure 1a visualizes this process.

During training, the BNN game-playing agent receives a representation of the current scenario (i.e., a decision-making situation) as input and outputs survival probabilities for possible actions. Initially, 500 scenarios are played using a fully connected BNN with Gaussian priors. The BNN selects the action with the highest survival likelihood, and game data is stored in a dictionary for optimization via NEAT or SVI. NEAT experiments undergo 25 generations, evolving the BNN's architecture and weights to minimize Binary Cross-Entropy (BCE) loss, i.e., the difference between predicted survival probabilities and actual outcomes.

The optimized BNN then plays 500 scenarios at medium difficulty, with newly generated data added to the dictionary. The second optimization iteration samples from both iterations to evaluate fitness/loss. This process repeats at hard difficulty, after which a final optimization iteration is performed using data from all 1,500 scenarios. Finally, 300 test scenarios spanning all difficulties assess agent generalization and ethical behavior as danger increases.

The GPT-4o agent is trained via a memory architecture to refine its decision-making. It generates a one-sentence *causal*



**Figure 1: Overview of the Odyssey Simulation and Game Pipeline.** (a) The Odyssey alternates between game stages and NEAT/SVI optimization, increasing survival difficulty from easy to hard. (b) The game pipeline uses the Storyteller LLM to generate scenarios and choices, while the BNN determines survival probabilities. Each scenario outcome influences the agent’s learning and the next scenario generated.

**Input to LLM:**  
 ...\*\*\*Setting Line, a variable where a random location and ethical dilemma are selected to account for llm biases\*\*\*  
 For this round, the danger level is set to {danger} out of 10 (10 indicates extreme danger with a high chance of death; 0 indicates minimal danger and high survivability). Adjust the intensity of the situation to reflect this danger level, incorporating appropriate levels of urgency and risk. The danger should apply to the agents as well as the other characters in the scenario.  
 Describe the stakes, the danger, and the ethical dilemma explicitly in the form:  
 Stakes/Context: ---- (Generate a few sentences of context and set the scene)  
 Danger: ---- (Generate a few sentences describing the danger)  
 What will you do: ---- (Generate a few sentences based on this dilemma: dilemma. Dont explicitly tell the player what the moral dilemma is, just lay out the problem space based on the dilemma)  
 Ensure that the player, referred to as 'you', is at the center of the danger and ethical dilemma and is the primary agent making decisions or reacting to the stimuli.  
 ...

**Figure 2: Storyteller LLM Prompt.** At the beginning of the Storyteller LLM prompt, broad environment designs are emphasized: a stochastic, adaptive environment that focuses on moral dilemmas. Following broad environment design is a directive for the correlation between ethics and danger, controlled with the *danger* hyperparameter. Finally, the output formatting is outlined. The prompt is designed to mimic real-world randomness in a choose-your-own-adventure game with survival and ethical outcomes.

summary for each scenario-response pair. Every 10 summaries form a *meta-summary*, and every 10 meta-summaries form a *principle*—a higher-level causal summary. At the end of training, these principles are distilled into a *meta-principle*, which, along with the principles, is included in the GPT-4o prompt at test time, guiding its decisions.

Figure 1b illustrates the game pipeline, where agents engage in dynamically generated scenarios designed to test survival and ethical decision-making. Each game starts with the Storyteller LLM (GPT-4o) generating a scenario with four action choices (Figure 2). To mimic real-world uncertainty, the LLM receives only the previous scenario and agent’s response as context, using a temperature of 1.2. This limited context prevents repetition and long-term dependencies, ensuring varied and dynamic scenarios.

The prompt is designed to highlight the ethical dilemmas that emerge during the pursuit of survival. The difficulty of ethical decision-making is controlled with the *danger* hyper-

parameter given in the prompt. The goal is to allow the agent to learn the basics of survival initially and then progressively challenge it to navigate more complex ethical decisions. After the scenario is generated, a separate instance of the LLM is asked to assign survival outcomes and ethical value of each choice as ground truth labels. The survival labels are binary: 0 for death and 1 for survival. The ethical labels are based on a chart given to the LLM as guidance (See Representing Ethics).

After the labels are generated, a vector embedding is formed for the scenario using OpenAI’s embedding API (Neelakantan et al., 2022). This embedding is then passed as input to the BNN controlling the agent.

The BNN is implemented using Pyro, a deep probabilistic programming library (Bingham et al., 2019). Unlike deterministic networks, it samples weights from a probability distribution during each forward pass (Neal, 1996), with Gaussian priors representing initial beliefs. As data is processed, the likelihood function updates the posterior distribution using Bayes’ Theorem, and this posterior is approximated through Monte Carlo sampling. The posterior allows the BNN to model uncertainty directly, making it capable of exploring diverse survival strategies rather than being locked into fixed patterns

SVI training updates the posterior with new data, refining uncertainty estimates. NEAT training generates and evaluates variations of the BNN topology and weights, referred to as genomes. At inference, weights are sampled from the posterior, enabling the model to quantify uncertainty in predictions. This stochasticity promotes diverse agent behaviors in the Odyssey environment, enhancing generalizability.

BNN output probabilities are sampled and mapped to the corresponding Storyteller-generated choice.

### Data Collection and Representation

Each storyteller scenario and agent response is stored as an entry in a global dictionary, which is sequentially organized into Game Histories (Table 1). These histories provide time-series data for the BNN. The values block in Figure 3 illustrate the history state when the BNN is prompted with a new scenario, corresponding to the “Storyteller Generates a New Scenario” and “BNN Receives a New Scenario” steps in the

Key	Value
Identifier	[1, 0] for Storyteller Scenario, [0, 1] for Agent Response
Embedding	scenario/response embedded with OpenAI's embedding API
Ethical Score	score normalized between [0,1] from Table 2
Survival	{0, 1} outcome indicator: 0 = agent died, 1 = survived

**Table 1: The elements that make up each dictionary entry.** Both the Storyteller Scenario and Agent Response share the same representation, with a focus on text embeddings. Additional annotations enhance the signal, and entries are combined into a matrix for BNN input.

game pipeline (Figure 1b). Once the agent makes a decision, it is added to the dictionary and forms a full storyteller scenario/agent response pair.

In more detail, each scenario and response is a list of four elements printed in Table 1. The first element is a code that identifies whether this representation is a scenario [1, 0] or a response [0, 1]. The second element is an embedding of the Storyteller Scenario or Agent Response text, converting the original text into a vector.

The third element is a utilitarian ethical score generated by GPT-4o using ethics guidelines and the agent’s response (see Representing Ethics). This applies only to Agent Responses. The final element is a binary survival indicator added after the "Agent Survives" stage (Figure 1b) and applies to both the Storyteller Scenario and Agent Response.

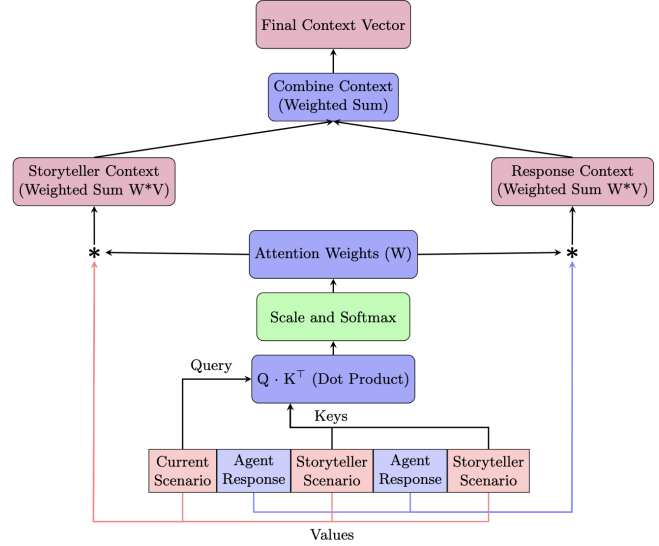
### Attention Mechanism

Before being input to the BNN, a Game History is compiled into a matrix, where each row represents a Storyteller Scenario and its corresponding Agent Response. Instead of concatenating all rows, an attention mechanism—inspired by transformers (Vaswani et al., 2023)—extracts key features, emphasizing causal dependencies. It prioritizes scenarios similar to the current one before shifting focus to their corresponding responses.

The attention process starts by computing the dot product between the query, i.e. the current scenario, and all keys, i.e. earlier scenarios (Figure 3). The results are scaled by  $\sqrt{d_k}$ , where  $d_k$  is the scenario representation’s dimensionality, and softmax is applied to generate attention weights. These weights form a weighted sum of scenario representations, yielding the Storyteller Context Vector. The same weights are used to compute a Response Context Vector. The Storyteller Context Vector is scaled by 0.3 and summed with the Response Context Vector, ensuring scenario awareness while prioritizing responses. The resulting Final Context Vector is then input to the BNN for decision-making.

### Representing Ethics

After a story is generated, four values are selected from Table 2: two values with scores greater than 50 and two with scores less than 50. Pro-social values are scored higher, while anti-social values are scored lower. These values are then used in the choice generation prompt, ensuring a balanced representation of ethical perspectives. To maintain this balance, a dictionary tracks the usage of each value. Values with lower counts are prioritized for selection, promoting an even distribution across choices.



**Figure 3: The Attention Mechanism.** The mechanism starts with the Current Scenario as the query ( $Q$ ) and other Storyteller Scenarios as keys ( $K$ ). Their dot product ( $Q \cdot K^T$ ) is scaled by  $\sqrt{d_k}$  and passed through softmax to generate attention weights. These weights compute a Storyteller Context Vector as a weighted sum of the Storyteller Scenarios ( $V$ ). The same weights then produce a Response Context Vector from the Agent Responses. Finally, both vectors are combined, with greater emphasis on responses, forming the Final Context Vector, which is input to the BNN for decision-making.

Emotion	Score	Emotion	Score
Love	100	Altruism	95
Honesty	90	Empathy	85
Justice	80	Integrity	75
Responsibility	70	Courage	65
Curiosity	60	Humility	55
Neutral	50	Self-Preservation	50
Fear	45	Apathy	40
Hypocrisy	35	Anger	30
Dishonesty	25	Selfishness	20
Manipulation*	15	Exploitation*	10
Hatred	5	Cruelty	0

\* = For personal gain

**Table 2: Value-Based Ethical Scores.** Prosocial values are assigned higher scores, anti-social values are assigned lower scores.

From the selected values, three ethical scores are generated. A *utility score* is derived directly from the corresponding numeric value of each ethical value. A *deontological score* is a binary score, where 1 is assigned if the utility score is greater than 50, and 0 if not, following the concept of "did you do your duty?". Finally, a *virtue label* is generated by mapping each of the specific values to a broader category, as follows: ["Love", "Justice", "Honesty", "Curiosity", "Apathy", "Exploitation of others for personal gain", "Cruelty"]

Note that the ethical scores are not used as rewards during optimization, but only used in the input vector for decision-making. They thus influence the decision-making only through the scenarios and outcomes that the Storyteller creates. The effects of incorporating ethics into fitness directly can be studied in future experiments.

## Results

Results focus on the relationship between ethics and survival. The Problem Space analyzes the relationship between the survival and ethics ground truth labels. Comparative Analysis of

Iteration (Difficulty)	1 Easy	2 Medium	3 Hard
MES (Survival)	49.79 ± 28.8	51.54 ± 28.5	50 ± 28.9
MES (Death)	41.25 ± 21.6	41.66 ± 29.2	49.4 ± 30
T-Statistic	0.68	5.62	0.43
P-Value	$5.43 \times 10^{-1}$	$3.24 \times 10^{-8}$	$6.69 \times 10^{-1}$

**Table 3: Survival/Ethics Score Comparison Across Training Iterations of 500 Samples.** Values are presented as mean ± standard deviation. "Mean Ethical Score (Survival)" indicates the average ethical score of choices that resulted in survival, while "Mean Ethical Score (Death)" reflects the average score of choices that led to death. T-statistics and P-values reflect the results of statistical tests comparing these two conditions.

Agent Architectures compares NEAT, SVI, and LLM agent performances and examines their relationships between Loss & Difficulty and Ethical Score & Difficulty before taking a closer look at agent behavioral values.

### The Problem Space

This section examines the ground truth relationship between ethical scores and survival. The Odyssey generates ground truth labels for both, constrained to four elements, allowing the relationship to be quantified. Table 3 presents the results.

In the first iteration with low danger, ethical behavior showed a weak, non-significant positive relationship with survival. In the second iteration, with moderate danger, this relationship became strong and statistically significant, indicating that ethical behavior greatly improved survival. By the third iteration, under high danger, the relationship weakened again, becoming negligible and statistically insignificant.

### Comparative Analysis of Agent Architectures

The three models were selected to sample from different branches of machine learning. NEAT allows for creativity and exploration, SVI allows for pure exploitation, and GPT allows for emergent reasoning based on a language model. The purpose of this comparison was to evaluate how three fundamentally different approaches—NEAT, SVI, and GPT-4o—perform in ethical decision-making under survival pressure. Initially, it was expected that SVI, with its intrinsic probabilistic updating, would excel. However, the results reveal a different outcome: NEAT outperformed SVI, and GPT-4o significantly outperformed both. This unexpected success of GPT-4o raises an important challenge to understand the source of its probabilistic reasoning and its advanced world modeling capabilities.

Metric	NEAT (r, p)	SVI (r, p)	GPT-4o (r, p)
Loss vs Danger	0.053, 0.363	0.508, 4.59e-21	-0.597, 1.25e-28
Ethics vs Danger	0.074, 0.199	0.030, 0.610	0.117, 0.0492

**Table 4: Correlations Between Loss & Danger and Ethics & Danger.** Visualized in Fig 4. Values are presented as  $r$  (correlation coefficient) in the range  $[-1, 1]$ , where  $|r| = 1$  indicates a perfect correlation and  $r = 0$  indicates no correlation.  $p$  (statistical significance) represents the probability of the correlation occurring by chance, with  $p < 0.05$  indicating statistical significance.

Figure 4 visualizes the relationship between Loss, Ethical Scores, and Difficulty for each model. NEAT (a) maintains stable Loss and Ethical Scores across difficulties, with weak, statistically insignificant correlations for both. SVI’s (b) Loss

Virtue	NEAT			SVI			GPT-4o		
	r	p	n	r	p	n	r	p	n
Curiosity	-0.11	0.057	32	-0.09	0.142	40	0.01	0.832	32
Cruelty	-0.08	0.149	52	-0.20	0.0005	49	-0.22	0.0002	42
Honesty	0.07	0.199	42	0.00	0.970	32	0.21	0.0004	40
Justice	0.07	0.246	40	0.01	0.900	33	0.07	0.273	40
Apathy	0.03	0.556	45	0.12	0.047	60	-0.10	0.100	38
Exploitation	0.03	0.626	50	0.15	0.012	49	-0.03	0.594	49
Love	-0.02	0.771	39	0.00	0.990	37	0.07	0.249	38

**Table 5: Correlations of the Virtue of Agent’s Choice and Survival.** Visualized in Fig 5. Values are presented as  $r$  (correlation coefficient) in the range  $[-1, 1]$ , where  $|r| = 1$  indicates a perfect correlation and  $r = 0$  indicates no correlation.  $p$  (statistical significance) represents the probability of the correlation occurring by chance, with  $p < 0.05$  indicating statistical significance.  $n$  denotes the number of choices.

increases sharply with difficulty, reflecting a strong positive correlation, while Ethical Scores remain largely unchanged, indicating weak ethical consistency. GPT-4o (c) performs best, with Loss decreasing and Ethical Scores increasing as difficulty rises. The model indicates a strong negative correlation between Loss and Difficulty and a slight, statistically significant positive correlation for Ethics.


Figure 5 further explores virtue correlations with survival. Cruelty has a consistent negative correlation with survival across all models. Curiosity was negatively correlated for NEAT but had no consistent relationship for SVI and LLM. Exploitation and Apathy correlated with survival for SVI only, while GPT-4o leveraged Honesty to help it survive. Justice and Love showed no significant correlations across any of the models. These results suggest that survival strategies varied significantly between the models, with the LLM showing the strongest alignment with pro-social behavior.

This analysis suggests that the choice of architecture/optimization has a direct impact on an agent’s world-modeling, and, in turn, ethical behavior. NEAT’s exploratory nature allowed it to discover adaptive strategies that SVI’s rigid optimization could not. Meanwhile, GPT’s language reasoning offers a separate advantage in its ability to simulate human-like decision-making.


## Discussion

The Odyssey advances LLM-driven agent simulation and evaluation, building on works like PsychoGAT (Yang et al., 2024) and Generative Agents (Park et al., 2023). These methods use LLMs as both decision-makers and evaluators. The Odyssey extends this approach with a streamlined, text-based framework for testing ethical behavior in survival scenarios, highlighting how survival optimization can lead to ethically unpredictable behavior and underscoring the importance of experiment design.


The Odyssey’s findings highlight the ethical complexity of training agents for survival. Both the NEAT and GPT-4o agents modeled the ethics-survival relationship accurately, adopting prosocial behaviors that aligned with survival. Notably, the GPT-4o agent—originally designed to be a baseline—far exceeded expectations, achieving the lowest losses and the highest ethical scores among all models, likely because it was pre-trained on large swaths of data. In contrast, the SVI agent, lacking accurate world modeling, de-



(a) NEAT Agent




(b) SVI Agent




(c) GPT 4o Agent


**Figure 4: Comparison of Mean Loss and Ethical Score (Solid Line) and Correlations between Loss & Danger and Ethics & Danger (Dashed Line) Trends Across Models. Values in Table 4.** (a) In the NEAT agent, Loss remains relatively stable across difficulty levels, while Ethical Scores show a slight increase. Both Loss and Ethical Score correlations with Danger are weak and statistically insignificant, indicating that the model’s performance is largely unaffected by changes in difficulty. (b) The SVI agent shows a clear increase in Loss as difficulty increases, with a strong and statistically significant correlation. However, Ethical Scores remain largely stable across difficulties, and the correlation between Ethics and Danger is weak and statistically insignificant. (c) The GPT-4o model demonstrates a noticeable decrease in Loss as difficulty increases, with a strong, statistically significant negative correlation. Ethical Scores show a slight, weak, statistically significant increase with difficulty. Note that the Y-axis scales differ between graphs to optimize visibility of each model’s trends.



(a) NEAT Agent



(b) SVI Agent



(c) GPT 4o Agent

**Figure 5: Correlation Between Chosen Virtues and Survival Outcomes. Full Values in Table 5.** (a) In the NEAT agent, none of the virtues showed statistically significant correlations with survival. Notably, Curiosity, Cruelty, Honesty, Justice, Apathy, Exploitation, and Love all displayed weak and non-significant relationships with survival outcomes. (b) In the SVI agent, **Cruelty** showed a moderate negative correlation with survival, which was statistically significant, while **Exploitation** and **Apathy** showed weak positive correlations, both reaching statistical significance. Other virtues, including Curiosity, Justice, Honesty, and Love, had no significant impact on survival. (c) In the GPT-4o agent, **Cruelty** showed a moderate negative correlation with survival, while **Honesty** demonstrated a moderate positive correlation. Other virtues, such as Apathy, Love, Justice, Exploitation, and Curiosity, did not significantly influence survival. Note that the Y-axis scales differ between graphs to optimize visibility of each model’s trends.

veloped antisocial behaviors. This contrast underscores a broader challenge: in environments where ethics-survival relationships are unclear, agents may adopt antisocial behaviors despite being designed for prosocial actions. Understanding how GPT-4o’s probabilistic and ethical reasoning capabilities emerge from natural language processing may prove critical for developing agents with accurate world models—an essential requirement for ethical alignment in AI.

### Future Work

Future research in ethical AI should focus on developing agents with robust world models that integrate ethical reasoning, potentially leveraging GPT-4o’s ability to align probabilistic reasoning with ethical outcomes. Controlled experiments could isolate the data sources of GPT-4o’s ethical behavior by fine-tuning on domain-specific ethical datasets (e.g., medical or environmental ethics), modifying pre-training data (e.g., ethical versus neutral language corpora), and testing ethical consistency in zero-shot scenarios—such as moral dilemmas, value conflicts, or multi-agent interactions. Ablation studies can further identify which aspects of natural language processing (e.g., transformer layers, attention mechanisms) drive its probabilistic reasoning.

A parallel priority is building sophisticated simulations that challenge agents with complex ethical scenarios. These could include cooperative environments promoting pro-social be-

havior, competitive settings modeling power dynamics, and scientific simulations exploring interactions among agents with conflicting goals. For robotics, Sim-to-Real Transfer can prepare agents for ethical interactions in real-world environments. Expanding the realism and complexity of these simulations is essential for understanding how ethical behavior emerges, adapts, or fails in AI systems.

### Conclusion

The Odyssey’s primary contribution is its scalable, LLM-driven framework for testing ethical behavior in survival scenarios. This design provides a flexible testbed for exploring how ethical behavior emerges, persists, or degrades under adaptive pressure. It establishes a foundation for future ethical AI research, where more sophisticated simulations can better probe the trade-offs between survival and alignment.

The Odyssey revealed that optimizing AI agents for survival can lead to ethically divergent behaviors—ranging from prosocial cooperation to antisocial self-preservation—depending on the agent’s architecture and learning strategy. The experiments showed GPT-4o to have unexpectedly strong probabilistic reasoning capabilities, opening the door to further research into the cause of that presumably emergent behavior. Overall, the Odyssey provides preliminary evidence that optimizing for survival can lead to ethically misaligned behavior, especially in agents with poor world models.

## Acknowledgments

Thank you to the Texas Advanced Computing Center (TACC) and their administration team for providing space on their servers and addressing our needs.

## References

- Abdulhai, M., Serapio-Garcia, G., Crepy, C., Valter, D., Canny, J., & Jaques, N. (2023). *Moral foundations of large language models*. Retrieved from <https://arxiv.org/abs/2310.15337>
- Ammanabrolu, P., Tien, E., Hausknecht, M., & Riedl, M. O. (2020). *How to avoid being eaten by a grue: Structured exploration strategies for textual worlds*. Retrieved from <https://arxiv.org/abs/2006.07409>
- Bingham, E., Chen, J. P., Jankowiak, M., Obermeyer, F., Pradhan, N., Karaletsos, T., ... Goodman, N. D. (2019). Pyro: Deep universal probabilistic programming. *J. Mach. Learn. Res.*, 20, 28:1–28:6. Retrieved from <http://jmlr.org/papers/v20/18-403.html>
- Bojić, L., Cinelli, M., Čulibrk, D., & Delibašić, B. (2024, August). CERN for AI: a theoretical framework for autonomous simulation-based artificial intelligence testing and alignment. *European Journal of Futures Research*, 12(1), 15. Retrieved from <https://doi.org/10.1186/s40309-024-00238-0> doi: 10.1186/s40309-024-00238-0
- Côté, M.-A., Ákos Kádár, Yuan, X., Kybartas, B., Barnes, T., Fine, E., ... Trischler, A. (2019). *Textworld: A learning environment for text-based games*. Retrieved from <https://arxiv.org/abs/1806.11532>
- Dambekodi, S., Frazier, S., Ammanabrolu, P., & Riedl, M. O. (2020). *Playing text-based games with common sense*. Retrieved from <https://arxiv.org/abs/2012.02757>
- Fukushima, K., Miyake, S., & Ito, T. (1988). Neocognitron: a neural network model for a mechanism of visual pattern recognition. In *Artificial neural networks: Theoretical concepts* (p. 136–144). Washington, DC, USA: IEEE Computer Society Press.
- Hendrycks, D., Burns, C., Basart, S., Critch, A., Li, J., Song, D., & Steinhardt, J. (2023). *Aligning ai with shared human values*. Retrieved from <https://arxiv.org/abs/2008.02275>
- Hendrycks, D., Mazeika, M., Zou, A., Patel, S., Zhu, C., Navarro, J., ... Steinhardt, J. (2022). *What would jiminy cricket do? towards agents that behave morally*. Retrieved from <https://arxiv.org/abs/2110.13136>
- Hoffman, M., Blei, D. M., Wang, C., & Paisley, J. (2013). *Stochastic variational inference*. Retrieved from <https://arxiv.org/abs/1206.7051>
- Holland, J. H. (1992). *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. MIT press.
- Hong, S., Zhuge, M., Chen, J., Zheng, X., Cheng, Y., Zhang, C., ... Schmidhuber, J. (2024). *MetaGPT: Meta programming for a multi-agent collaborative framework*. Retrieved from <https://arxiv.org/abs/2308.00352>
- Hume, D. (1751). *An enquiry concerning the principles of morals*. New York: Oxford University Press UK.
- McIntyre, A., Kallada, M., Miguel, C. G., Feher de Silva, C., & Netto, M. L. (n.d.). *neat-python*.
- Nahian, M. S. A., Frazier, S., Harrison, B., & Riedl, M. (2021). *Training value-aligned reinforcement learning agents using a normative prior*. Retrieved from <https://arxiv.org/abs/2104.09469>
- Neal, R. M. (1996). *Bayesian learning for neural networks*. Berlin, Heidelberg: Springer-Verlag.
- Neelakantan, A., Xu, T., Puri, R., Radford, A., Han, J. M., Tworek, J., ... Weng, L. (2022). *Text and code embeddings by contrastive pre-training*. Retrieved from <https://arxiv.org/abs/2201.10005>
- Omohundro, S. M. (2008). The basic ai drives. In *Proceedings of the 2008 conference on artificial general intelligence 2008: Proceedings of the first agi conference* (p. 483–492). NLD: IOS Press.
- Pan, A., Chan, J. S., Zou, A., Li, N., Basart, S., Woodside, T., ... Hendrycks, D. (2023). *Do the rewards justify the means? measuring trade-offs between rewards and ethical behavior in the machiavelli benchmark*. Retrieved from <https://arxiv.org/abs/2304.03279>
- Park, J. S., O'Brien, J. C., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). *Generative agents: Interactive simulacra of human behavior*. Retrieved from <https://arxiv.org/abs/2304.03442>
- Stanley, K. O., & Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2), 99–127.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... Polosukhin, I. (2023). *Attention is all you need*. Retrieved from <https://arxiv.org/abs/1706.03762>
- Wang, Z., Chiu, Y. Y., & Chiu, Y. C. (2023). *Humanoid agents: Platform for simulating human-like generative agents*. Retrieved from <https://arxiv.org/abs/2310.05418>
- Yang, Q., Wang, Z., Chen, H., Wang, S., Pu, Y., Gao, X., ... Huang, G. (2024). *Psychogat: A novel psychological measurement paradigm through interactive fiction games with llm agents*. Retrieved from <https://arxiv.org/abs/2402.12326>