MORTAL AGENTS WORKSHOP - ALIFE 2025

FADING FACES

WHEN AGENTS FORGET WHO YOU ARE

David Freire-Obregón

SIANI, Universidad de Las Palmas de Gran Canaria, Spain

INTRODUCTION

- **Perceptual forgetting** differs fundamentally from **erasure-based forgetting**. While the latter implies deletion of stored information, the former involves declining ability to recognize what was once familiar.
- This research draws inspiration from human **cognitive aging**, where individuals may retain symbolic identity knowledge but lose perceptual recognition capabilities.
- **Recognition** is a foundational type of social memory. For both humans and artificial agents, knowing "who" depends on perceptual, symbolic, and embodied processes.
- Understanding how perception deteriorates in agents provides insights into **cognitive mortality**, memory systems, and the nature of identity recognition.

RESEARCH QUESTION

PERCEPTUALFORGETTING

COGNITIVE PROCESSES

MAIN QUESTION

How does the progressive degradation of visual fidelity in facial representations affect the recognition capabilities of artificial agents, particularly when symbolic identity remains unchanged?

This research explores the boundary between perceptual access and identity recognition, questioning whether agents' memory works more like human recall than conventional data storage.

HYPOTHESIS

We hypothesize that recognition accuracy will decline precipitously as visual fidelity decreases, even while symbolic identity remains intact and accessible to the agent.

THEORETICAL BASIS

This work draws from cognitive theories of mortality where forgetting is not binary but dissociative - one may recall a name but not a face, or recall an emotion but not the individual.

RESEARCH IMPLICATIONS

By demonstrating this symbolic-perceptual dissociation, we explore how artificial systems may develop human-like memory limitations through perceptual constraints rather than data loss.

METHODOLOGY

S I M U L A T I O N

MESA FRAMEWORK

Agent-based simulation implemented using the Mesa framework, where agents move randomly within a 2D toroidal grid. Upon encountering neighbors, agents attempt to recognize each other using personal neural networks for face identification.

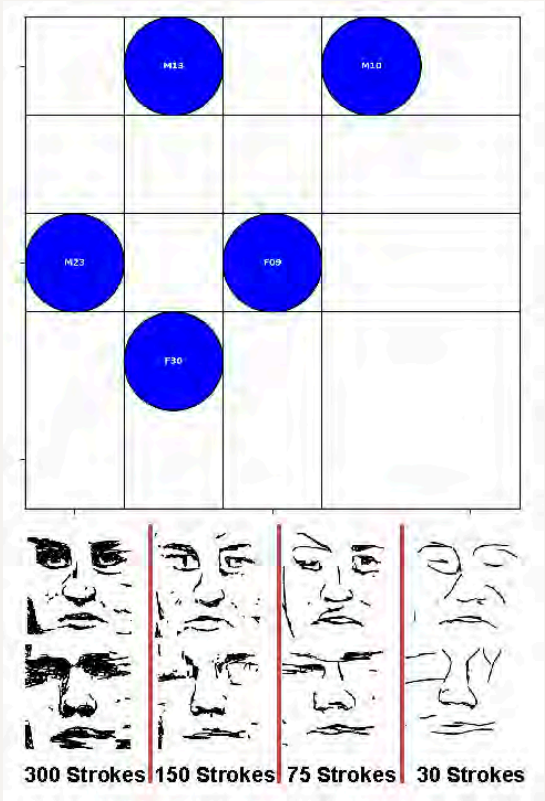

D A T A S E T S

KDEF & CLIPASSO

Facial images from the KDEF dataset provide neutral expression identities. Visual degradation is implemented using CLIPasso sketches at four abstraction levels: 300, 150, 75, and 30 strokes—modeling progressive perception decline.

// NEURAL ARCHITECTURE

TECHNICAL

NEURAL NETWORK

RESNET18 CNN

Each agent is equipped with its own ResNet18 CNN for face recognition. During the first 70 simulation steps, agents learn by encountering others and incrementally updating their classification layer to recognize neighboring faces.

Close-up of an Abstract Line Painting with Faces — Designecologist / Pexels

PERCEPTUAL DECAY

VISUAL DEGRADATION

Faces gradually deteriorate from detailed 300-stroke renderings to increasingly abstract versions with 150, 75, and finally 30 strokes. This progressive degradation simulates perceptual decline while agent identity remains unchanged.

VISUAL DEGRADATION

SIMULATION DETAILS

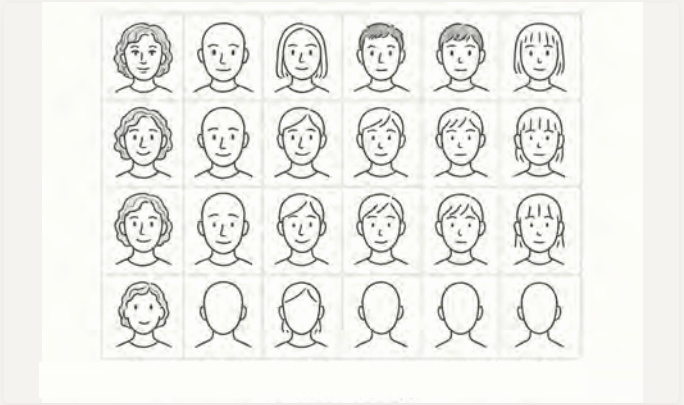

SIMULATION PROCESS

AGENT ENVIRONMENT SETUP

The experiment implements a 5×5 toroidal grid where agents randomly move and interact with neighbors. Each agent has a unique facial identity from the KDEF dataset, which is rendered at varying levels of abstraction using CLIPasso sketches.

When two agents enter neighboring cells, they attempt mutual recognition, testing their neural networks' ability to correctly classify increasingly degraded visual stimuli.

PROGRESSIVE FACE ABSTRACTION

Faces are rendered with four decreasing levels of visual fidelity using the CLIPasso sketch technique:

300 STROKES

150 STROKES

75 STROKES

30 STROKES

RESULTS

- Initial **peak recognition accuracy** observed near step 70, with agents successfully consolidating recognition capabilities during the learning phase with high-fidelity images (300-stroke sketches).
- As visual quality deteriorates (**150→75→30 strokes**), a **precipitous drop** in recognition confidence occurs, reaching less than **35% accuracy** by step 130.
- Symbolic identity (unique agent IDs) remains intact throughout the simulation, while visual recognition capabilities severely degrade.
- Results confirm the hypothesis: agents "forget" not through erasure of internal representations, but through **perceptual mismatch** between stored representations and increasingly abstract visual inputs.

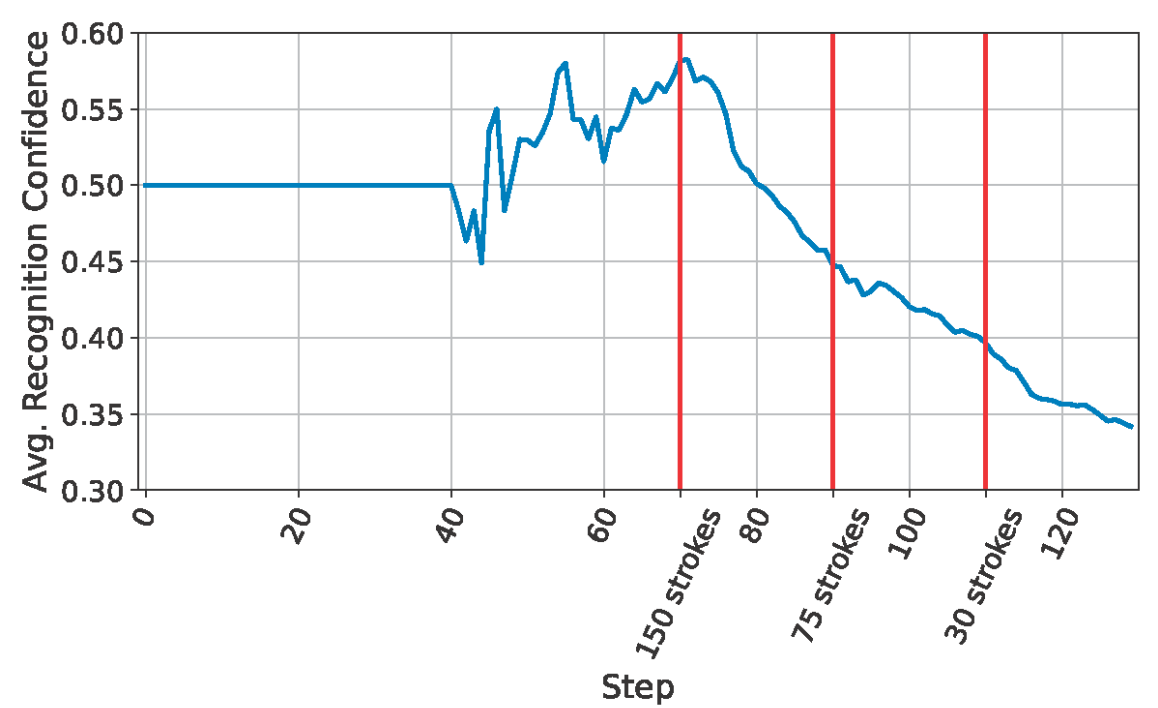

Figure 2: Average recognition confidence over time. Red lines indicate degradation steps (150, 75, 30 strokes).

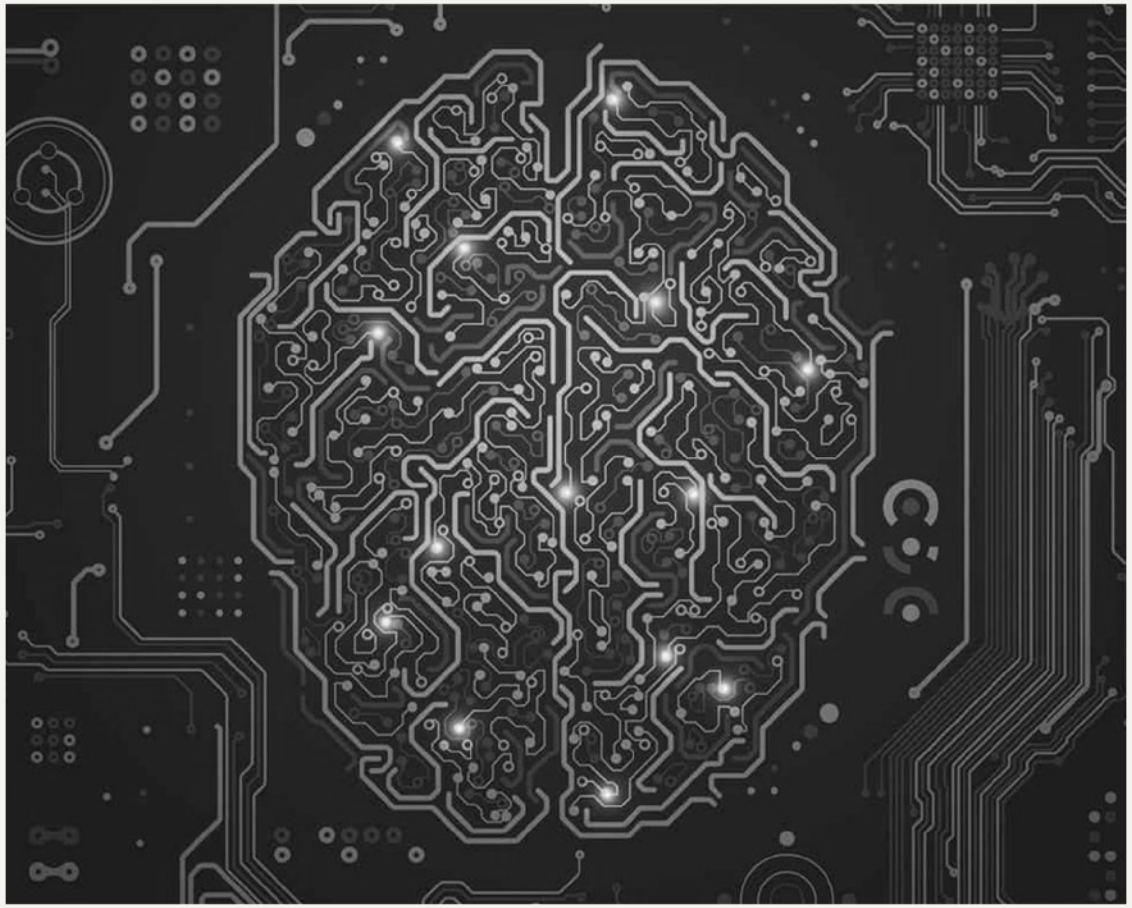

KEY INSIGHTS

- **Symbolic-Perceptual Dissociation** emerges as agents retain symbolic identity (unforgettable unique ID) while perceptual access to that identity decays, mirroring human cognitive phenomena where one can recall a name but not a face.
- **Artificial Life Implications:** Forgetting manifests not as erasure from memory but as a mismatch between internal representations and external inputs, suggesting a structural understanding of forgetting in artificial systems.
- **Embodiment:** The research connects perception, recognition and embodied cognition, demonstrating how cognitive decline affects social interaction between artificial agents within their environment.
- **Cognitive Mortality:** Results suggest that mortality in AI systems can be understood not merely as functional decay, but as erosion of familiarity and recognition — a form of deteriorating cognitive relationship with the world.

CONCLUSIONS

- This research provides a **symbolic account** of memory and forgetting in embodied agents through face-based social interaction.
- **Forgetting emerges** not from erasure of memory but from **perceptual mismatch** between internal representations and corrupted inputs.

FUTURE DIRECTIONS

- Exploring the **relationship between symbolic and perceptual layers** in more complex agent architectures.
- Investigating **cognitive mortality** as a lens for understanding embodiment and abstraction in artificial memory systems.
- Developing frameworks for **enhancing perceptual resilience** in social agents facing imperfect sensory input.

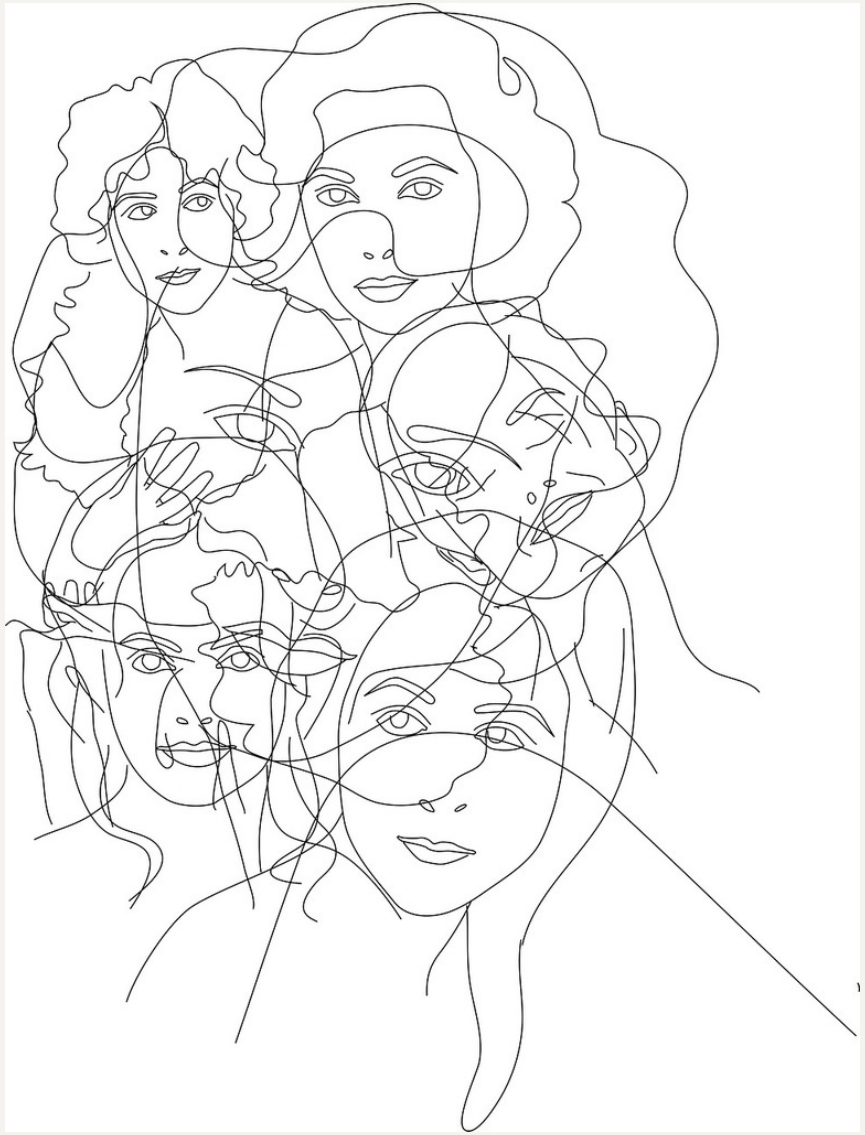

Image: Women Face Drawing (ed6313), Public Domain — Picryl

THANK YOU