Morphological Murmurations

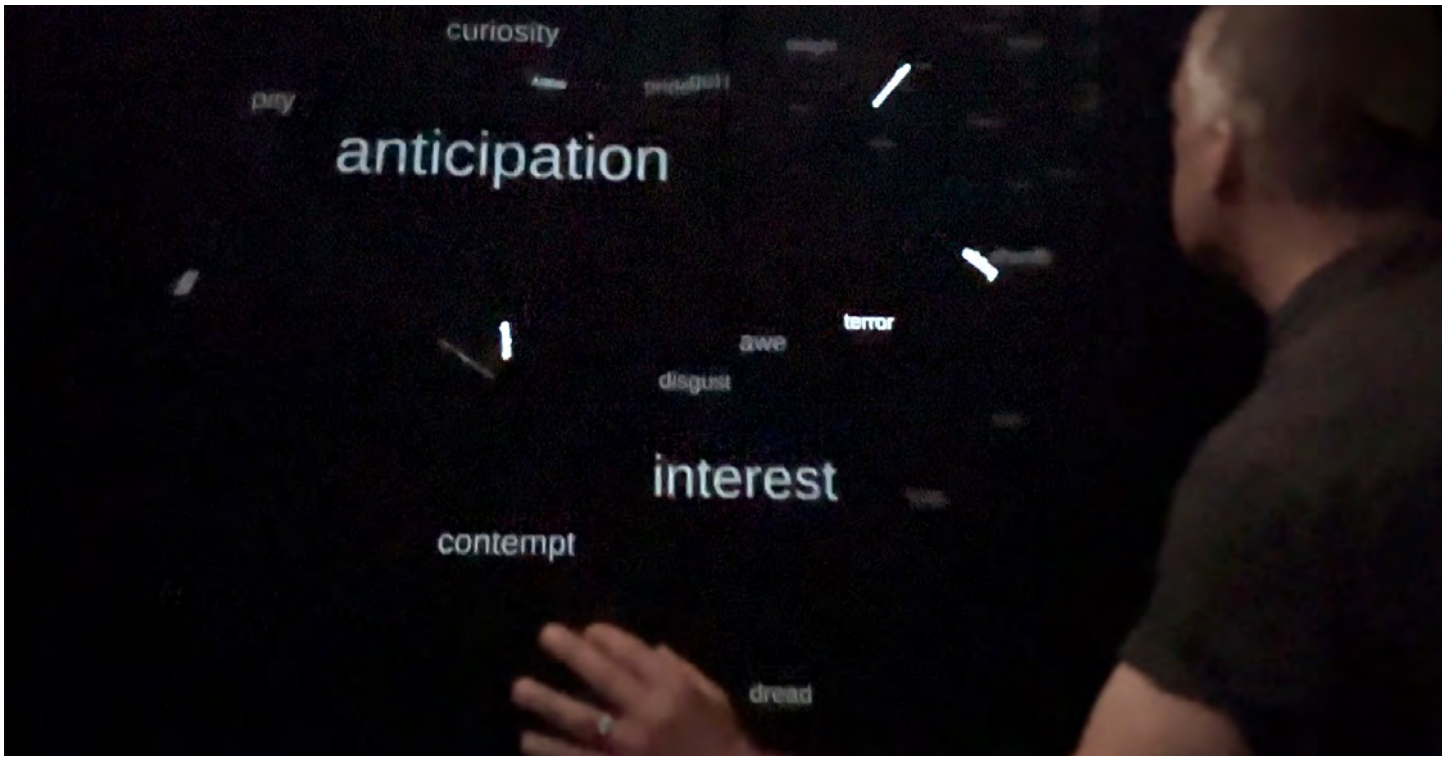

Theodore Koterwas
tkoterwa@ed.ac.uk
University of Edinburgh, Edinburgh, United Kingdom

Morphological Murmurations engages its audience kinaesthetically with language and draws on the flocking behaviour of birds to broaden consideration of how AI conceptualises language and the roles it might play in both facilitating and limiting the creation of meaning. Utilising a depth camera, game engine software, deep reinforcement learning (DRL) and voice cloning text-to-speech (TTS), the installation places visitors within a virtual 3D representation of the embedding space of Google BERT, a large language model (LLM) where a flock of artificial agents responds to their movements, visually and sonically activating words through a synthetic clone of the artist's voice. This is language assembled not through computational logic, but through patterns of bodies working together: that of the visitor and those of the flock. Through the intuition of movement, muscle memory and spatial awareness, the visitor explores the environment of a putative "other mind" (the LLM) through an alternative model of how embodied minds (birds) coordinate movement and create meaning with their bodies. Evaluation of the logs of interactions during a public exhibition of the work suggests the types of meaning produced and compares these with how an LLM assesses the meaningfulness of word sequences.

Keywords: LLM, Biomimicry, Reinforcement Learning, Distributed Cognition.

DOI 10.34626/2026_xcoax_027

Introduction

Morphological Murmurations was exhibited during the Scottish AI Alliance event *AI Can't Write Symphonies and Neither Can You*, 28-9 June 2025 in Glasgow, United Kingdom. In response to the provocation in the event's title, the installation proposes creative expressions like symphonies are artefacts of something more essential to what it means to be human: the ability to experience and shape thought and expression through movement and embodied interactions with others in physical and conceptual space. It explores differences between human minds and algorithmic models in meaning-making and promotes a more inclusive consideration of ways of thinking and communicating than the computational view of cognition underlying the LLMs currently dominating the implementation of and discourse around Artificial Intelligence.

Fig. 1. Installation view of *Morphological Murmurations*. Online video documentation: (https://theodorekoterwas.com/morphological-murmurations)

Background

The idea for the work originated when I intuitively wrote "trust pets" in response to a question during a *Writing the Wrongs of AI* workshop run by Pip Thornton (Thornton 2024). One need only consider guide dogs, carrier pigeons, and other "working" animals to recognise the high degree of confidence we place in their response to unpredictable situations. From millennia of observing and training animal behaviours, we recognise our own propensity to act predictably without requiring what Anthony Chemero refers to as "mental gymnastics:" the conscious manipulation of mental representations, symbols or computation (Chemero 2009). We also see in these animal behaviours something beautiful, such as in Starling murmuration, or terrifying, such as in the swarming of bees. These behaviours signify meaning through the coordination of moving bodies. *Morphological Murmurations* attempts to mimic flocking behaviour to propose an alternative to the logic of the LLM in navigating the semantic structure of language.

Word Embeddings as the Linguistic Environment of Computational Minds

The word embeddings of an LLM are the “environment” in which it operates. Early construction of embeddings relied on an approach to language, Latent Semantic Analysis, that is inherently spatial, presuming an underlying semantic structure based on proximity: words surrounded by similar words are themselves similar (Dumais et al. 1988). In the case of the Transformer architecture of modern LLMs, this structure is transformed through layers of computational manipulation using Queries, Keys and Values to determine the relevance of specific tokens (word and sub-word units) in a particular context (Vaswani et al. 2017). While it may be working in incomprehensible dimensions, these dimensions encode spatial assumptions about the nature of language.

To comprehend the environment in which an LLM operates, we need to view the embeddings in a more familiar dimensional space. A three-dimensional representation is created by reducing the high dimensional vectors of each word to the three most distinctive values and then mapping them as Cartesian coordinates. This process reveals spatial semantic relationships, clustering words and, famously, exposing parallel directional links between conceptually related pairs of words such as “king” and “male” and “queen” and “female” (Mikolov et al. 2013). Because they were derived from the use of language by the authors of the training material, these relationships in the embeddings can be seen as a shared semantic space of the authors, which in the case of Google BERT, is hundreds of thousands of contributors to the English Wikipedia and BookCorpus datasets (Devlin et al. 2019; Wu et al. 2016; Zhu et al. 2015). The idea that machine generated spaces represent meaningful structures or the contents of human mental processes has been proposed, explicated and contested in works of art and theoretical papers, particularly those examining the “latent” spaces of visual models (Schaerf 2024; 2025). The embedding space represents something more foundational. If the latent space is a space of possibilities and hallucinations, the embedding space is the environment these models “move”, homing in on the most probable, and therefore, according to a computational theory of mind, most meaningful word.

Distributed Cognition and Embodied Language

The computational theory of mind is contested by theories of distributed cognition which hold that a significant amount of human behaviour does not require mental representations or the manipulation of symbols. Examples such as catching a ball, walking upright, even playing instruments can be accounted for through dynamic systems analysis of bodily capabilities and environment affordances (Varela et al. 2017; Chemero 2009; Pfeifer et al. 2007). Language appears to defy these accounts: what are words if not mental representations, signifiers or symbols? Andy Clark (2009) presents a middle ground in which language scaffolds extended cognition: words are a means of labelling and there-

fore grouping things in our environment to reduce higher order problems, they allow us to rehearse actions, control selective attention and anchor thought and reason, and they appear to us as objects themselves. In addition, we utilise our bodies as we form ideas into words. Learned and habitual movements shape our thoughts as much as subtle actions of "body language" might belie them. We move our hands, arms, head, torso and even our hips and feet to shape concepts, grasp at relationships, and add weight and rhythm (McNeill 2005; Arnheim and McNeill 1994; Cassell et al. 2000).

An embodied view of cognition considers behaviour as emerging from a dynamic system comprised of the mind, body and the environment (De Bruin et al. 2021). Such a view inspires an investigation of how meaning might be produced through the navigation of an environment comprised of spatially arranged signifiers. A Phenomenological view of language holds that "it is the lateral relation of one sign to another which makes each of them significant, so that meaning appears only at the intersection of and as it were in the interval between words" (Merleau-Ponty 1978, 42). These intersections and intervals require not just existing spatial relationships between words, but new connections made between them through acting in the world. It is proposed in this installation these connections are made as much by the body as the mind and through the coordination of multiple bodies.

Flocking and Swarming as Alternative Exploratory Method

Flocking and swarming are forms of emergent, collectively intelligent behaviour that require no consideration of mental representation or symbolic thought to describe or simulate. They are modelled algorithmically by calculating the position of agents using three parameters: alignment with other agents, separation from them and cohesion of the group while following a target, often a "leader" agent or prey. These parameters are not representative of "thoughts" or "calculations" but purely of spatial relationships between agents. Navigating word embeddings according to these parameters offers an alternative to the way LLMs operate on them and explores whether the meaning we find in the spatial patterns of group coordination translates to the spatial composition of language. Exploring language through biomimicry inherits from experiments in language development by autonomous agents and the Synthetic Method of cognitive research employed in Artificial Life (Pfeifer et al. 2007; Steels 2012). While Artificial Life approaches take the dynamic, situated and socially-embedded nature of intelligent behaviour as a premise, they have been challenged for their underlying assumption that the dynamics of these systems represent a reality that is predictable (Birhane 2021). Utilising biomimicry as an exploratory method in art accepts and utilises this unpredictability to question theories of cognition through the empirical method of artistic practice.

Fig. 2. Detail view of flocking agents navigating word embeddings.

Evaluation and Reflection

Often artworks are evaluated through visitor feedback. In this case we are expressly interested in alternative experiences of language where meaning is implicit, created through evocative bodily and spatial navigation. By examining sequences of words triggered in the logs we can evaluate their potential to evoke meaning in an associative way and how this compares to a purposeful attempt to navigate concepts. We can supplement this analysis with observations made of visitors interacting with the work and verbal feedback they provided.

Evaluation of Logs

The software logged every word triggered by the flocking agents. A comparison of two logs can provide insights into the effectiveness of the installation in producing meaningful sequences of words. The “audience” log recorded public interactions during the second day of the exhibition. Prior to opening I interacted with the installation for filming. Because I was purposefully attempting to demonstrate meaningful interactions, the “artist” log of these can help gauge the effectiveness of the installation in providing similarly meaningful interactions for the public.

The logs were analysed programmatically to sample sequences of words triggered. Reading over these sequences, it is difficult to compare them for meaningfulness. They contain a mix of strings of disjointed concepts (“worth puzzle mutual importance consistency parody satirical”), lists of semantically related words (“fatigue longing annoyance dread terror embarrassment terror”), and moments of poetic potential (“weighted fantasy gauge secret nightmare dreams truth”). The human desire to find patterns and relationships leads to potentially oversubscribing meaning to these sequences through association and intuition. While this is part of the point of the work, it presents an obstacle to a more objective assessment. The intention of the work is to propose an alternative method of navigating the embedding space than the one

used by LLMs, so using an LLM to evaluate the results offers a means of analysis. Submitting the sequences for evaluation by GPT-5.2 through the University of Edinburgh's ELM service (University of Edinburgh, n.d.) reveals different concepts of "meaning" created through the computational approach to language modelled by LLMs and the more associative, embodied approach of a human.

Some examples of these evaluations are "These lines carry relatively coherent human/psychological or situational meaning compared to the more technical/repetitive ones" and "Strong emotional content + social circumstances + cognition/epistemic states." For GPT-5.2 meaning appears to be correlated with concepts such as "human/psychological", "emotional states", "social/political concepts", "evaluative/interpretive" and "situational" rather than "technical", "statistical", "procedural" or "repetitive." GPT-5.2 also identified "clusters" of meaning: "A coherent 'life outcomes & adversity' cluster", "Clear 'stress/fear/discomfort reactions' cluster" and "Strong, coherent adversity/outcome grouping." Concepts such as "coherence", "clarity", "consistency" and "strength" are correlated with meaning. This analysis demonstrates the spatial character of the concept of meaning in LLMs: meaning is derived from words being conceptually grouped in a clear and coherent manner, particularly words associated with human characteristics rather than technical or statistical ones. By this analysis, spatially navigating words with one's body makes sense.

Observations and Feedback

During the second day of the exhibition I dedicated a few hours to systematically observing audiences interacting. My observations were primarily intended to assess how effectively the work encouraged active bodily engagement, and I would only approach visitors if they seemed to struggle to interact. I found that of the sixteen individuals or groups, five required encouragement to approach the work and all but one began moving their bodies with purpose once they sensed movement of the words corresponding to their approach. The variety and dynamism of movement varied from walking around to crouching and waving arms. Four visitors required an explanation of how the interactivity worked, but then explored the words with intent, moving from cluster to cluster. A few visitors engaged me in further discussion. One visitor found the experience of navigating words with their body "therapeutic", another described the experience as "visceral" and "moving," and a third relayed that "there was a moment when the word 'violence' flew by and [they] freaked out a bit." One particularly engaged visitor, having navigated effectively, discussed their prior reading around spatial relationships across languages and animal communication. The installation provoked interactions within groups. In one family, the child danced and grabbed for the birds, and the parents took turns navigating with intent and remarked, "That's really cool, really good." One of a pair of visitors described how they compared words: her interaction produced the words "punishment" and "threat" while her companion produced more posi-

tive terms, resulting in feelings of being judged even though she knew it couldn't judge her. The other artist exhibiting in the room said she felt judged when the piece kept repeating "boring" as she described her own work to a visitor even though she knew it couldn't hear her.

Conclusion

Morphological Murmurations explores alternative processes of constructing meaning from a spatial model of language. By placing the visitor in the embedding space of a LLM accompanied by flocking agents the installation proposes language is enacted through a dynamic system of bodies interacting in a linguistic environment. A comparison of human and machine evaluations of the resulting sequences reveals different concepts of meaning, the human being more associative and intuitive, and the machine based on semantic clustering and distinctions between "human"-related characteristics and "technical" properties.

Acknowledgements. The work was informed by conversations with Jane McKie and Stuart King. Jane, a poet, and I shared a fascination with Starling murmuration and the way meaning emerges in language. Stuart, a mathematician, is particularly interested in fluid dynamical processes.

Link: https://theodorekoterwas.com/morphological-murmurations

References

Arnheim, Rudolf, and David McNeill.
1994. "Hand and Mind: What Gestures Reveal about Thought." *Leonardo* 27 (4): 358. https://doi.org/10.2307/1576015

Birhane, Abeba.
2021. 'The Impossibility of Automating Ambiguity'. *Artificial Life* 27 (1): 44–61. https://doi.org/10.1162/ARTL_A_00336

Cassell, Justine, Matthew Stone, and Hao Yan.
2000. "Coordination and Context-Dependence in the Generation of Embodied Conversation." In *Proceedings of the First International Conference on Natural Language Generation*, vol. 14, 171–78. https://doi.org/10.3115/1118253.1118277

Chemero, Anthony.
2009. *Radical Embodied Cognitive Science*. A Bradford Book. MIT Press.

Clark, Andy.
2009. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford University Press.

De Bruin, Leon, Gunnar Declerck, Michael David Kirchhoff, and Elmarie Venter.
2021. "Toward an Embodied, Embedded Predictive Processing Account." *Frontiers in Psychology* 12: 543076. https://doi.org/10.3389/fpsyg.2021.543076

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova.
2019. "BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding." In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Volume 1 (Long and Short Papers), 4171–86. Minneapolis: Association for Computational Linguistics.

Dumais, Susan T., George W. Furnas, Thomas K. Landauer, Scott C. Deerwester, and Richard A. Harshman.
1988. "Using Latent Semantic Analysis to Improve Access to Textual Information." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '88), 281–85. Washington, DC: ACM. https://doi.org/10.1145/57167.57214

McNeill, David.
2005. *Gesture and Thought*. University of Chicago Press.

Merleau-Ponty, Maurice.
1978. *Signs*. Northwestern University Studies in Phenomenology and Existential Philosophy, no. 6. Northwestern University Press.

Mikolov, Tomas, Wen-tau Yih, and Geoffrey Zweig.
2013. "Linguistic Regularities in Continuous Space Word Representations." In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (Atlanta, Georgia), June, 746–51.

Pfeifer, Rolf, Josh Bongard, and Simon Grand.
2007. *How the Body Shapes the Way We Think: A New View of Intelligence*. MIT Press.

Schaerf, Ludovica.
2024. "Reflections on Disentanglement and the Latent Space." *arXiv*. arXiv:2410.09094. https://doi.org/10.48550/arXiv.2410.09094

Schaerf, Ludovica.
2025. "The Evolving Nature of Latent Spaces: From GANs to Diffusion." In *Ethics and Aesthetics of Artificial Intelligence Conference* (EA-AI'25), arXiv:2510.17383. https://doi.org/10.48550/arXiv.2510.17383

Steels, Luc, ed.
2012. *Experiments in Cultural Language Evolution*. Vol. 3. John Benjamins Publishing Company. https://doi.org/10.1075/ais.3.04ste.

Thornton, Pip.
2024. "Writing the Wrongs of AI." *Pip Thornton* (blog). Accessed 10 February 2026. https://pipthornton.com/2024/08/12/writing-the-wrongs-of-ai/

University of Edinburgh.
n.d. "ELM." Accessed 15 February 2026. https://elm.edina.ac.uk/elm/elm.

Varela, Francisco J., Evan Thompson, and Eleanor Rosch.
2017. *The Embodied Mind: Cognitive Science and Human Experience*. 2nd ed. The MIT Press. https://doi.org/10.7551/mitpress/9780262529365.001.0001

Vaswani, Ashish, Noam Shazeer, Niki Parmar, et al.
2017. "Attention Is All You Need." In *Advances in Neural Information Processing Systems 30* (NIPS 2017), 5998–6008. https://doi.org/10.48550/arxiv.1706.03762

Wu, Yonghui, Mike Schuster, Zhifeng Chen, et al.
2016. "Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation". arXiv, October 8, 2016. arXiv:1609.08144. https://doi.org/10.48550/arXiv.1609.08144

Zhu, Yukun, Ryan Kiros, Rich Zemel, et al.
2015. "Aligning Books and Movies: Towards Story-Like Visual Explanations by Watching Movies and Reading Books." In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)*, December, 19–27. https://doi.org/10.1109/ICCV.2015.11