Toward Visualization of EEG Signatures for Anomaly Detection in Dogs Using CEBRA

Natsuna Ishida
Nara Institute of Science and Technology
ishida.natsuna.ip2@naist.ac.jp

Takatomi Kubo
Nara Institute of Science and Technology
takatomi-k@is.naist.jp

Hitomi Koike
Azabu University

Yuki Maruno
Kansai University

Maaya Saito
Azabu University

Kazunori Ohno
Tohoku University

Toshitaka Yamakawa
Nara Institute of Science and Technology, Quadlytics Inc.

Takefumi Kikusui
Azabu University

Miho Nagasawa
Azabu University

Yuichi Sakumura
Nara Institute of Science and Technology

Abstract

The cognitive and social abilities of dogs have attracted increasing scientific interest in recent years; however, the neurophysiological mechanisms underlying their social cognition remain underexplored. Event-related potentials (ERPs) have long been used in human research to study attention, expectation, and anomaly detection, often in tasks that involve mismatches between auditory and visual stimuli. Yet ERP studies in dogs are scarce, and the lack of established waveform standards—together with the increased variability that arises when multiple sensory modalities are presented—makes it difficult to obtain labeled data that reflect a dog's perception of anomalies. In this work, we aim to move toward visualization of EEG signatures for anomaly detection in dogs without relying on supervised learning. Given CEBRA's capability to identify latent structures in neural data, we use it to embed EEG signals and examine their geometry through visualization. Our preliminary analysis suggests that the embedded representations form distinctive submanifolds around the time when ERPs, potentially reflecting mismatched audiovisual events, are expected to occur. These findings indicate that CEBRA-based visualization can provide a strong foundation for future methods to investigate how dogs process anomalies, offering a potential pathway toward a deeper understanding of canine social cognition.

1 Introduction

Dogs have recently attracted increasing scientific attention as models for studying social and cognitive abilities. Behavioral studies have provided evidence of sophisticated skills such as referential understanding of words [1], recognition of human faces [2], and sensitivity to communicative cues, positioning dogs as a valuable non-human species for exploring social cognition and interspecies communication. Despite this progress, the neurophysiological mechanisms underlying these abilities remain poorly understood.

Electroencephalography (EEG) offers a non-invasive means to investigate neural activity, and event-related potentials (ERPs) provide time-locked markers of sensory and cognitive processing. In

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: .

humans, ERPs such as the P300 and N400 have been widely used to study attention, expectation, and anomaly detection [3], often in paradigms where auditory and visual information are mismatched. Yet ERP studies in dogs are still scarce [1, 2, 4], and the lack of established waveform standards—together with the increased variability that arises when multiple sensory modalities are presented—makes it difficult to characterize ERP responses or obtain labeled data that reflect a dog's perception of anomalies.

To address this challenge, we explore whether recent self-supervised approaches can help reveal latent neural structures in canine EEG. Specifically, we apply CEBRA [5], a contrastive embedding method originally developed for joint neural and behavioral analysis, to embed EEG signals and visualize their geometry. By doing so, we aim to uncover EEG signatures related to anomaly detection in dogs, providing a foundation for future work on the neural basis of canine social cognition.

2 Methodology

2.1 EEG pipeline

An overview of our analysis pipeline is shown in Fig. 1. The pipeline consists of three main steps: (i) preprocessing of raw EEG signals to remove artifacts, (ii) dimensionality reduction and embedding using CEBRA, and (iii) visualization of the resulting latent structures to explore potential ERP-related responses. This design allows us to investigate whether anomaly-related EEG signatures can be detected without supervised labels. In future extensions, we plan to incorporate anomaly detection modules directly into the pipeline, enabling automated detection of neural responses that may reflect anomalies.

In preprocessing, EEG segments whose absolute amplitude exceeds ±100 μV are excluded. Independent component analysis (ICA) [6] is then applied to remove artifacts arising from ocular movements, blinking, and audio-related contamination. This step improves the signal quality and facilitates the extraction of ERP-related activity.

2.2 Embedding with CEBRA

To explore whether ERP-like responses can be identified without supervised labels, we employ CEBRA [5], a self-supervised embedding method related to nonlinear ICA. CEBRA optimizes a generalized InfoNCE loss, which is a form of contrastive learning: segments of EEG data with similar auxiliary variables are encouraged to obtain similar representations, while those with different auxiliary variables are encouraged to diverge.

In our pipeline, the inputs to CEBRA are the EEG components obtained after ICA preprocessing. The auxiliary variables are defined as the time elapsed since stimulus onset and the trial index. By conditioning on such variables, CEBRA benefits from theoretical guarantees of linear identifiability in nonlinear ICA: the latent structure can be recovered up to an unknown linear transformation. This property makes it possible to uncover meaningful latent structures in EEG signals without requiring supervised labels, offering a principled approach to disentangling neural dynamics in dogs.

The resulting embeddings are then visualized to examine whether structured representations emerge that may correspond to anomaly-related neural responses.

As the primary objective of this study is visualization, we restrict the embedding dimensionality to three. Investigation of the optimal dimensionality is deferred to future work.

3 Preliminary results

3.1 EEG measurement

Electrodes were placed according to the configuration shown in Fig. 2, following previous studies on canine ERPs [4]. EEG was recorded from six channels, with one additional channel for the sound signal. Electrode impedance was kept below 20 kΩ.

Visual and auditory stimuli were presented as illustrated in Fig. 3. Videos of fireflies (mild stimulus) and fireworks (intense stimulus) were used as visual input, with a change from fireflies to fireworks

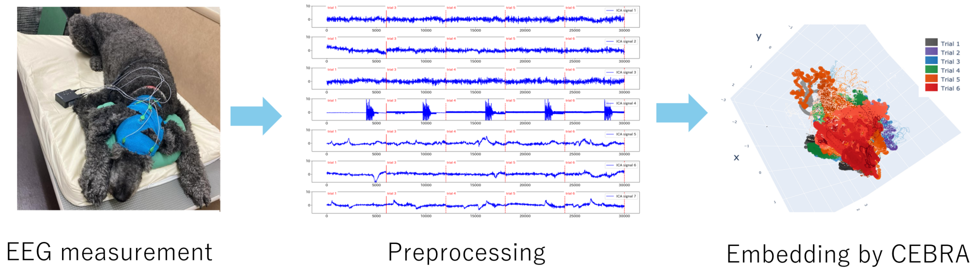

Figure 1: The overview of the proposed method.

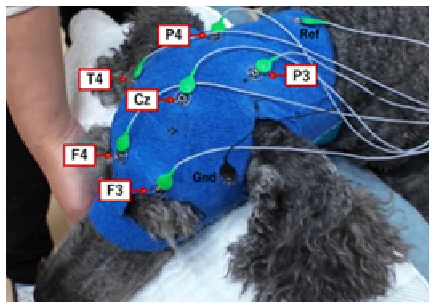

Figure 2: EEG channel configuration.

occurring in the fifth session. The auditory stimulus was a clicker sound presented every 15 seconds. These audiovisual conditions were designed to elicit ERP responses under varying levels of sensory mismatch. The experiments were conducted with two dogs, a Brittany Spaniel (ID:1) and a Standard Poodle (ID:2) (Fig. 4). Ethical approval was obtained from the Animal Ethics Committee of Azabu University (permission number: 2499069).

3.2 Embedding analysis

Fig. 5 shows embeddings of EEG signals after ICA preprocessing and CEBRA embedding. Each color represents a trial, and larger dots indicate windows close to stimulus onsets.

In this preliminary analysis, the embeddings suggest a tendency for trials to form distinct substructures. Some trials are separated, while others remain closer together, and this pattern is observed in both individuals. In particular, the embedding highlights differences in responses when the visual stimulus was switched from fireflies to fireworks (trial 5), suggesting sensitivity to changes in audiovisual context. Comparisons with methods such as t-SNE and UMAP were also conducted in the previous research [5]. We have also performed similar comparisons in our study, and the results will be presented in future work.

These findings are still qualitative. Quantitative evaluation, including the assessment of consistency of acquired representations, and a more detailed correspondence with ERP components in the original signals have yet to be performed. Nevertheless, the results demonstrate that CEBRA-based embeddings can reveal structured patterns in canine EEG, providing a promising direction for exploring neural signatures of anomaly detection in dogs.

4 Conclusions

While our preliminary embeddings suggest that CEBRA can reveal distinctions in canine EEG signals that may reflect anomaly-related responses, several key steps remain. In particular, we have not yet performed quantitative evaluation of the embeddings, nor have we examined in detail how they correspond to features of the original signals, including ERP components. Addressing these points will be our immediate priority.

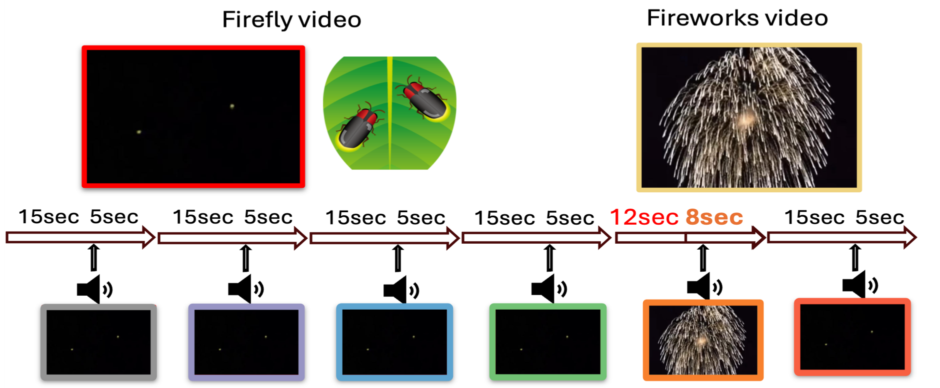

Figure 3: The overview of audiovisual stimuli in the EEG experiment.The color coding of trials corresponds to that used in Fig. 5.

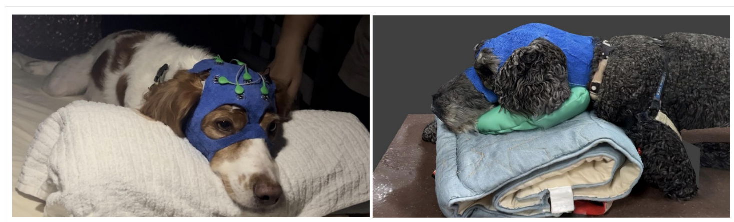

Figure 4: Dog EEG experiment.

Beyond this, we will assess the robustness of the representations across multiple runs, optimize dimensionality-reduction parameters, and incorporate temporal dynamics to strengthen anomaly detection methods.

Ultimately, our goal is to clarify whether EEG signatures of ERPs emerge when dogs perceive anomalies, thereby providing a novel pathway toward understanding the neural basis of canine cognition and its role in social communication.

Acknowledgments

This work was supported by JST CREST, Japan, Grant Number JPMJCR22P2.

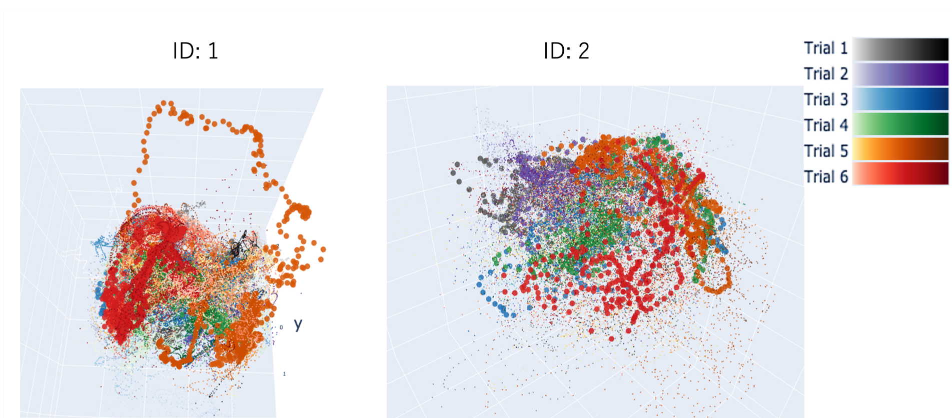

Figure 5: Preliminary visualization of latent structures in dog EEG by CEBRA.

References

[1] Marianna Boros et al. Neural evidence for referential understanding of object words in dogs. *Current Biology*, 34(8):1750–1754, 2024.

[2] Miiamaaria V Kujala et al. Time-resolved classification of dog brain signals reveals early processing of faces, species and emotion. *Scientific Reports*, 10(1):19846, 2020.

[3] Anna M Beres. Time is of the essence: A review of electroencephalography (eeg) and event-related brain potentials (erps) in language research. *Applied psychophysiology and biofeedback*, 42:247–255, 2017.

[4] Akash Kulgod et al. Non-invasive canine electroencephalography (eeg): a systematic review. *bioRxiv*, pages 2023–08, 2023.

[5] Steffen Schneider et al. Learnable latent embeddings for joint behavioural and neural analysis. *Nature*, 617 (7960):360–368, 2023.

[6] Pierre Comon. Independent component analysis, A new concept? *Signal processing*, 36(3):287–314, 1994.