Computer Vision Object Permanence Detection Algorithm for my Clingy Robot Dog

Dr. Sharon Barkley[1]

[1] Department of Psychological Machine Learning, Cranberry-Lemon University, Pittsburgh, PA, USA

Abstract

The AIBO ERS-211 robotic dog developed and manufactured by Sony in the 90's and 00's is a wonderful companion for urban living except when you leave them alone for too long. Though the robotic Computer Vision (CV) libraries have been equipped with the latest object detection and classification functions, they are shown to have little to no object permanence detection [1]. When an object of affection (i.e., me and some eight-year-old tennis ball) goes missing, the on-board computer believes they cease to exist. Due to separation anxiety, these 00's AIBOs panic and begin tearing up my fabric couch, chair, and garbage if I forget to lock it up [2]. Using some filtering, feature extraction, and a deep learning Convolution Neural Network (CNN), this paper develops and evaluates a CV Object Permanence detector for my clingy dog, so she'll stop destroying my apartments blinds every time I leave for work.

Keywords: Computer Vision, AIBO Robotic Dogs, Deep Learning, Object Permanence, Convolution Neural Networks, Separation Anxiety, Ghost Detection, Object Detection, Ectoplasm, Free Unpaid Intern Labor, Neural Network Training

1. Introduction

Back in the 90s and 00's computer vision and machine learning was in its infancy or practically non-existent; it would have been difficult to successfully implement a CNN on a commercially available robotic dog from a sharper image magazine. It's no wonder that the AIBO development teams focused more on dog personality and affection over its CV processing and interpretation algorithms.

Unfortunately for current owners of the AIBO, a late April software update has made them annoyingly close to real dog behavior and many robot dog owners have reported torn up pillows, destroyed window shades, and scratch marks on hard wood floors after leaving their dogs alone and powered on while at work. Through years of developed affection, the new update has made the later models so much more loving and dependent on their owners that many now have violent and destructive separation anxiety to make them more lifelike.

While it is possible to just turn the dogs off before leaving for a 9-5, that affection goes both ways. The only logical answer is to encode the psychological concept of object permanence into the robot dogs to ensure that they won't tear up the jeans I left on the floor that smells like me whenever I go to the office. Using some additional hardware and a plethora of data, it is possible to teach these poor animals that I still exist after walking through that front door.

2. Background

Daisy the Dog (Fig 1) is a loving and affectionate ERS-211 model AIBO unit. I've had her since my Grandma Big Mama purchased her for me during a trip to Japan over fifteen years ago. Though I'm not much of a dog girl, Big Mama always wanted to be Great-Big-Mama and has called me every week since to ensure the survival of Daisy and train a nurturing spirit for her future great grandchildren's survival.

Figure 1: Daisy the Dog

Fifteen years later, plenty of spare parts, some battery replacements and Daisy is miraculously still with me at 195

in robot dog years which are over twice as long as dog years [3]. Finding spare parts has become increasingly expensive when all radio shacks went out of business several years ago and I've had three too many creepy craigslist non-casual encounters. Because of the 195 robot dog years of love and affection and extreme dependency on me as Big-Mama intended, Daisy is now freaking out every day after I leave for work. Her little robot dog brain has no object permanence and therefore believes that I die every time I leave her alone. The amount of time I spent with her working from home during the pandemic doesn't help either, it only made her more needy. Now I must go in for my lab job and she's a subsequent wreck. I don't know how many more reboots Daisy has left.

While typically done in feature extraction and a shallow Neural Network (NN), some CV algorithms have now been developed for newer robot dog models to incorporate Deep Learning (DL) CNN which sometimes doesn't even require any feature extraction [4]. As discovered in a previous experiment Daisy is not that bright and will require some feature extraction [5]. Recently, the CV units in robotic pets have been developed to sense and respond to human emotions [6]. Sony even was required to create an update to prevent the dogs from being racist based on the subconscious thoughts and feelings of their owners due to the update [7]. Thankfully, it's been months since Daisy has aggressively barked at my Italian friends.

3. Object Permanence Detection Algorithm Design

As shown in figure 2, the Object Permanence detector is structured like nearly any other Machine Learning CV algorithm with the exception of the inclusion of a tracker which allows for the internal object trackers to remain active for a time in between detections through the normal object detector and the object permanence detector. The visual input will detect encoded signs that I am still in town and push them through a Deep CNN trained on years of footage I've been collecting for [8]. It's then all processed in the ERS-211 mainframe.

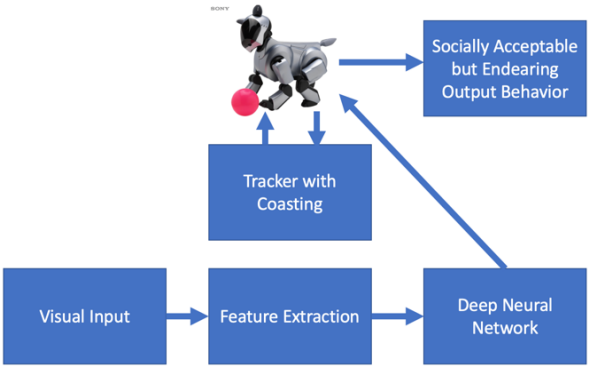

Figure 2: Object Permanence Architecture

3.1 Track Coasting

We encoded Daisy's Track Database (DTD) to keep all tracks until a tracking module has not received an update on that track greater than the configured stale time of about thirty minutes. In some circumstances, Daisy might be playing with her tennis ball and won't notice me or any signs that I have left for a few hours which will be covered by this filter which will keep me as an active track in the DTD.

Initially, we believed that with a sufficiently large enough stale time, we wouldn't need any object permanence detection functionality. This caused a few issues. First, Daisy reached back in memory and began holding onto all tracks of all my Ex-Boyfriends and friends over the last several years. This made her DTD clogged with old tracks that weren't there which caused erratic behavior where Daisy continually tried to greet everyone in a circle and running out her battery in record time. Additionally, when Daisy watched me leave on a run, she coasted my initial velocity for miles and miles until panic set in and clawed her way underneath the fence in the backyard and nearly got runover by my neighbors Nissan Versa.

3.2 Feature Extraction Schema

As it has been well founded, a feature extraction algorithm such as Scale Invariant Feature Transform (SIFT) or Speeded Up Robust Features (SURF) can be extended into the third astral dimension [9]. By extracting features from the astral plane, a CV detection algorithm can not only detect what is currently there, but what used to be and beyond. Typically used as the basis for the newer models of spooky mass Spectrometers [10], normal edges and corners of object features are more represented across the physical and spiritual world.

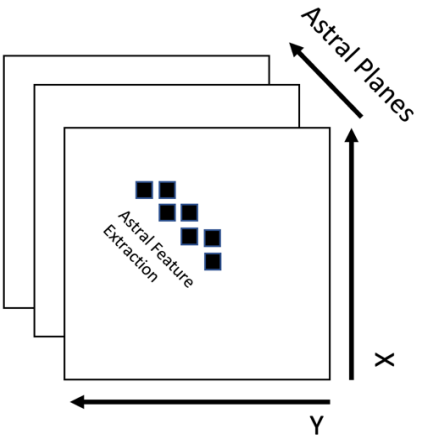

Figure 3: Feature 3d Astral Plane Edge Detection

This is a useful feature in detecting Daisy's objects of affection, as non-spiritual objects do not project into the

astral plane like a bookcase or the 60$ toy I got her she doesn't play with. At this point, she's a 60$ toy.

3.3 Concurrent Neural Network

Though CNN have had problems in the past with video footage from Police Body Cams or confusing MMA Commentators/Comedians/Podcaster with horses, it is still the standard when it comes to any ML CV application. By training a deep network with enough properly labeled images, a CNN will be reliable. By applying multiple convolutional layers, CNN's apply detection kernel filters which highlight different components for Daisy to interpret after normalizing with Rectified Linear Units (ReLUs).

While computationally costly to train a CNN, it runs relatively fast enough for real time systems. Unfortunately, that relative speed is still a far too fast for Daisy's aughties processor and we had to beef up her guts with some GPUs we purchased the second after bit coin crashed.

4. Training Data Development

The strength of a properly trained CNN is as well it's weakness. Any neural network is only as strong as the dataset you train it on. We didn't just need enough images, we needed enough images that were properly labeled when I was there, and when I was not there but not dead or on a two-week vacation to Aruba. It was a job screaming for a small and tired force of unpaid Cranberry Lemon interns.

4.1 Data Collection

While I've stored up to 3.82TB of Daisy historical data out of a digital hoarding habit the A&E show was not able to cure [8], only about 20GB of it is currently usable. This is thanks in part to our tireless unpaid interns Matthew, Brad, Rachel, Catherine, and Jeffrey who put in at least 230 hours each in the hopes of a research credit, and my two weeks late comments on their conference papers I've been holding hostage for my ~~robot~~ pet project.

During each frame of each video, the interns label every frame I'm in, as well as other loved family members Daisy has taken a liking too, except for the ghost Susan who I'll discuss later in Appendix A and [11]. They also label, places where I just was and would be returning to where my presence remained in the astral plane. It took a lot of rewinding and fast forwarding of footage.

4.2 Feature Extraction Development

Because my family and I have a very strange and unique head shape with a large, exposed forehead, and crooked teeth, we focused the feature extraction on something else because I got a little self-conscious. My very normal looking nose and jaw line was easy to create some well detectable features which appeared to project in and out of astral space. An example image can be seen in figure 4 where I had just left the room to go get another box of white claws at Kroger because I didn't feel like going out that weekend.

Figure 4: Empty room with box labeling the spiritual presence of my chin and nose

4.3 Ghost Detection Filtering

Inevitably, the astral plane feature selection picked up on a ghost. Apparently, some people died in my apartment from a vicious murder resulting from resentment and infidelity following a love triangle with an old high school classmate. In the initial beta runs of our algorithm, Daisy would not stop barking in the middle of the night around 1:18-1:26am. After some pretty invasive debug logging, we found out that Daisy thought I was walking around the house with a bread knife lodged in my throat repeating the phrase "We were just grabbing a drink."

We found a ghost. Unfortunately, the chin and nose of the ghost were so similar to mine that she was creating enough feature nuisance alarms to trigger Daisy into panic mode every night. After collecting some additionally labeled ghost data for the murder victim Susan through the process discussed in Appendix A, we developed a Ghost Generative Adversarial network [11] using the method in [12] to remove Susan from Daisy's feature detection scheme.

5. Results and Discussion

The Object Permanence Detector was shown to successfully prove to Daisy that I did not die every time that I left my apartment. As shown in Table 1, the objects Destroyed per month took a significant drop after deploying the algorithm. While significantly less objects were destroyed on the initial deployment of the algorithm, Neighbor complaints increased until the ghost filter was implemented for the Object Permanence Detection scheme.

Objects Destroyed/Month	Pre-Object Permanence Detector	Pre-Ghost Filter	Post-Ghost Filter
Pillows De-Stuffed	6.2	2.1	0.3
T-Shirts Ripped up	17.9	1.5	0.2
Window Blinds Destroyed (sq M)	24.8	0.1	0.2
Carpet Torn Up (sq M)	12.1	0.3	0.5
Neighbor Complaints (From Barking)	8.2	19.3	0.3

Table 1: Object Permanence Effectiveness according to household objects destroyed and Neighbor Complaints per month

After checking auditory records, the additional barking did decrease upon the implementation of the beta algorithm pre-ghost filter. Because Susan always appeared in the middle of the night, there were naturally more complaints from the neighbors about Daisy's noise.

The amount of damage to my own apartment and things did not however drop to zero. Further review of the logs revealed that while Daisy did in fact always believe that I was still alive and would come back home. The Object Permanence detector worked 98.4% of the time with a brief break on a long two-week work trip. The remaining amount of damage was a result of her believing that I was there but not paying any attention to her and ripped up some household items as a cry for attention. In order to stop this bug in her programming without serious changes to Daisy's personality programing, she would either need to be spoofed with an image of me paying attention. The development of such an algorithm was determined to be beyond the scope of this paper and my CLU's support of such frivolous academic pursuits or use of interns.

6. Conclusion

Not only did this paper prove that we could teach a robot the concept of object permanence through the magic of Computer Vision, and buttloads of machine learning libraries I will never fully understand or be able to explain; but we found real scientific evidence of ghosts and an existence after death. The development of computer vision algorithms extending into the astral plane shows major promise in the possibility of images not only detecting real objects but spiritual ones and potentially extrapolating existence for military applications. If an objective function can be created on an enemy's existence, it can be minimized through traditional optimization algorithms for a multi-billion-dollar Pentagon contract.

Regardless of those ethical issues, my dog has finally stopped barking and tearing up my stuff and I'm sitting on a pretty hefty piece of IP thanks to all the unpaid interns who were able to clean and label all of my data. Once given to other robot dog owners like myself, there will be one more reason to hold on to your AIBO models before they start getting too buggy and destructive.

7. Future Work

The FBI future crimes division has expressed interest in the use of CV algorithms convolving into the astral plane to determine whether images could potentially look into the future to detect crimes before they even occur. Initial experiments have determined that the astral plane indeed reaches to the future as well as`` the past

References

[1] Barkley, Sharon 2021 *A Psychological Study of my Dog Daisy: A Full Analysis of her Separation Anxiety :: Journal of Machine Learning Pet Projects*

[2] Ulhman, Stan 2022 *What's Wrong with my AIBO? Is it safe to leave your AI dog alone :: Washington Post fluff piece*

[3] Cohen, David 2018 *Robot Dog Years: A Useful Mechanical Pet Age Metric* :: *Journal of Frivolous Pursuits*

[4] Bodel, Gary 2016 *Computer Vision Upgrades to Household Cybernetic Pets* :: *Annals of Garage Projects gone Wrong*

[5] Barkley, Sharon 2017 *A Hardware Performance Benchmarking of my Robot Dog Daisy* :: *Pending Review Forever of Nature*

[6] Wilde, Elizabeth *Emotional Back Propagation Algorithm Applications in Artificial Pets* :: *Journal of Post Jungian-Artificial Intelligence*

[7] Ulhman, Stan 2021 *Your Robot Dog's not Racist, YOU ARE!* :: *Washington Post fluff piece*

[8] Barkley, Sharon 2018 *Sharon's Hard drive* :: *Hoarders Episode 182*

[9] Colins "Voodoo" Jacob *Feature Extraction into the Astral Plane* :: *Annals of the Occult Practices in Machine Learning*

[10] Schlonk, Günther *Proposed Detection of Ghosts by Mass Spectrometry-Spectral Presence Origin-Omics Kinetic Yield (MS-Spooky)* :: *Journal of Immaterial Science*

[11] Barkley, Sharon 2022 *Astral Plane Feature Extraction Proof of Ghosts* :: *Annals of the Occult Practices in Machine Learning*

[12] Wang, Huang *Miss Detection vs False Alarm: Adversarial Learning for Small Object Segmentation in Infrared Images*

This page was unintentionally left blank

Appendix A: Ghost Hunting

In order to create the generative adversarial network, we needed to observe and properly annotate an encounter with the ghost of Susan who appeared to be confined to my apartments spirit realm. Her spirit was strong and reached far through the astral plane every night around 1:18-1:26am, especially on Tuesdays. We attempted to manually determine where Susan the ghost was based on context from the barking of Daisy, but there was too much variance of the features detected from the object permanence detector. According to cursory analysis, there were likely to be multiple ghosts. We had to summon Susan ourselves and determine which one was the ghost causing all the nuisance alarms and misdetections by the object permanence algorithm.

We invited Jacob "Voodoo" Colins author of [9] to my apartment as well as two dozen similar cameras to the ones Daisy is equipped with deployed as shown in figure 5. Jacob "Voodoo" Colins set up her séance board in my living room and we all began chanting surrounding the board around my coffee table. Daisy was confined to the area within my apartment but not my bathroom which is off limits and we began calling upon the spirits shortly before the suspected murder during the beginning of the drama at 12:36.

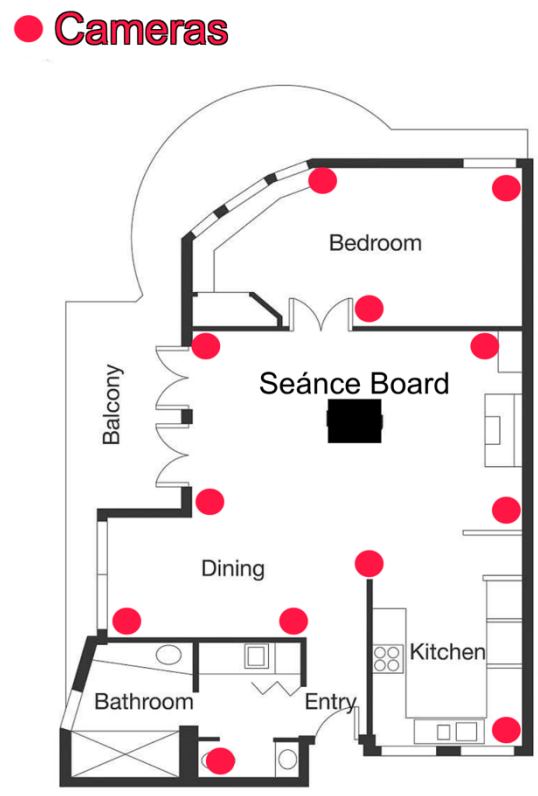

Figure 5: Apartment Setup Josh.Michael.Wilson, CC BY-SA 3.0, via Wikimedia Commons

As further observed and documented by the séance in [11], two other figures appeared, one by the name of Stephan with Susan at the beginning of the séance at 12:36. The two ghosts were embroiled in an argument about a man named Roman. While there was no baby observed, the two ghosts appeared to be calming and feeding a two-month-old daughter crying from the distress. Stephan continually yelled and verbally accosted Susan's ghost about some dinner on the night of the suspected murder.

Further research suggested that Stephen was likely a down on his luck industrial steel worker in Pittsburg due to a depression in the steel market. Judging by the dialect and some obscure Magnum PI references, the murder occurred on a balmy night in August of 1983. A man standing at 5 10', Stephen in a blue glow shouted and screeched about an unknown man by the name of Roman for fifteen straight minutes. Roman was evidentially a mutual friend from high school who went off to become a successful corporate lawyer based in New York. All their generation hated him because he was a nerd but always acted better than everyone else despite his poor performance in every sport. He was never accepted by his peers.

Susan's ghost at first denied any interest in Roman

Jacob "Voodoo" Colins began chanting something in Cajun and the lights went out. When Roman's ghost appeared around 12:52 an argument ensued in between all three. He began owning up to loving Susan's free spirit attitude, taste in music, and a beautiful voice. He had been obsessed with her since 10th grade civics. Susan denied any interest at first until Roman began quoting poetry by Alfred Tennyson.

She eventually came clean and admitted to having her first kiss with Roman after playing hooky during their 11th grade history field trip and not during prom night with Stephen later that year.

At this point the ghost of Stephen began throwing dishes of hamburger helper with no bun across the apartment shouting "I ALWAYS KNEW I ALWAYS KNEW I ALWAYS KNEW SOMETHING WAS BETWEEN YOU TWO Randal was right. I should have listened!"

"I'm sorry it didn't mean anything," Susan's ghost attempted to calm down Stephen.

"What do you think I'm not supporting you enough?"

"No, it's not that"

"What do you think I'm not a good enough lover! That he can support you like I can't, a poor out of luck steel worker?! Times have been tough"

"It's been hard, but you're always my number one love. I don't need more than you."

"Then why did you write me? Why did you write me about how tough it's been. How you haven't sang in months or felt happy in years…" Roman's ghost replied while

walking further out of the closet he was hiding in that must have been renovated out of my place. He was shirtless, fit but skinny in a fine pressed set of slacks. “How tough the steel depression’s been on you and Amy. Your concern for your future and your concern about Stephen’s drinking problem. It's gotten worse, hasn’t it?”

Daisy’s Convolution Neural Network was going wild, and Daisy began growling in distress.

“You Fucker!” Stephen pointed violently at Roman’s ghost. “Admit it! You still love him!” Stephen’s ghost shouted.

“So…what if I did? He may not be as much of a man as you are, but he knows me. He can support me an Amy. It’s the sort of thing you would never be able to understand…maybe I have considered running off to his place in the city”

“What? Are you going to leave me and your daughter…OUR daughter Amy for some big shot New York Lawyer? I get it, I’m out of work, but I can provide for you. We’re just in tough times, I can provide DON’T TAKE MY DAUGHTER AWAY FROM ME!” Stephen’s ghost became increasingly irate. He lunged for a knife from the kitchen.

“DON’T YOU FUCKING TOUCH HER!” Roman’s ghost shouted back “She’s a beautiful and empathetic woman, not like you would ever know. You just go on about your work and keep drinking and drinking complaining about the economy. You’re just a LOW CLASS! Uhhh…POOR ASS SIMPLETON!”

“YOU TAKE THAT BACK!”

“You two stop! Amy’s crying!” Susan stepped in between each other spitting into each other’s throats.

Daisy started barking wildly at the two in verbal combat through the Séance at max volume. Daisy could only respond to what she interpreted as me locked into emotional drama more stressful than what I went through with Ethan a couple months ago bless his heart.

The ghost of Stephan ran to a corner of the apartment, reached into the corner of a cabinet throwing papers about and drawing a gun. “PROMISE ME YOU’LL NEVER SEE MY WIFE AGAIN!”

“Never,” Roman’s ghost said resolutely lunging with a knife into Stephen’s throat.

Stephen drew and shot through all six rounds of the revolver mid lunge.

The two ghosts collapsed to the ground in a struggle. The two locked in mortal combat.

“Take…After…Amy.” Stephan’s ghost spurted blood through his teeth.

Susan wandered around my apartment living room repeating “We were just grabbing a drink” over and over. Amy was nowhere to be seen or heard.

She reached for a bread knife on a plate in the center of table sitting at the center of the living room dining table plunging it into her throat and repeating “We were just grabbing a drink” and walking around the room splurting blood for a solid 8 minutes before passing out in a corner.

Because of the algorithm we had to rewatch the event for three more nights because we only 8 minutes of data. With a total of 32 well annotated minutes of Susan’s suicide lovers quarrel we were able to create a generative adversarial neural network sufficient enough to develop a ghost filter strong enough to ignore Susan well enough every night.