

# Robot Synesthesia: A Sound and Emotion Guided Robot Painter

Vihaan Misra<sup>1</sup>, Peter Schaldenbrand<sup>1</sup>, and Jean Oh<sup>1</sup>

**Abstract**—If a picture paints a thousand words, sound may voice a million. While recent robotic painting and image synthesis methods have achieved progress in generating visuals from text inputs, the translation of sound into images is vastly unexplored. Generally, sound-based interfaces and sonic interactions have the potential to expand accessibility and control for the user and provide a means to convey complex emotions and the dynamic aspects of the real world. In this paper, we propose an approach for using sound and speech to guide a robotic painting process, known here as *robot synesthesia*. For general sound, we encode the simulated paintings and input sounds into the same latent space. For speech, we decouple speech audio into its transcribed text and the tone of the speech. Whereas we use the text to control the content, we estimate the emotions from the tone to guide the mood of the painting. Our approach has been fully integrated with FRIDA, a robotic painting framework, adding sound and speech to FRIDA’s existing input modalities such as text and style. In two surveys, participants were able to correctly guess the emotion or natural sound used to generate a given painting more than twice as likely as random chance. On our sound-guided image manipulation and music-guided paintings, we discuss the results qualitatively.

## I. INTRODUCTION

FRIDA [1], [2] is a collaborative robotic framework for creating real-world paintings. FRIDA can produce an artistic visualisation that is consistent with the human artist’s intent expressed in the text form. Whereas text description in language can succinctly express semantic concepts in an abstract manner, it can be challenging to fully express the complexity and nuances of sounds or emotions, e.g., people use emojis to express their emotions in textual communication.

Art has traditionally served as a powerful medium for expressing and evoking human feelings and emotions. Emotions form an integral part of the painting process, and it is crucial to understand the relationship between visual art and emotions to capture the painter’s actual intent. Studies find that different visual elements such as color, shape, and composition can evoke different emotions in viewers [3], [4]. By focusing on sounds and the emotions behind it, we aim to initiate a more nuanced perceptual understanding of the painting process, which, downstream, can also be applied to richer understanding of ordinary paintings. Additionally, a speech interface can improve the usability and accessibility of the system for users with motor, linguistic, and cognitive impairment [5].

Toward intuitive and powerful painting generation, we propose an approach to use the sound and speech input




Fig. 1: **Robot Synesthesia** connects a robot painter’s action space directly to user driven sonic interactions. For *Speech guidance*, sound is decoupled into language and emotion whereas with *Natural sound guidance*, the audio itself drives the content of the painting.

modalities for painting, introducing *Synesthesia FRIDA (S-FRIDA)*. Synesthesia is a perceptual phenomenon in which a person may perceive visuals when listening to sounds.

In contrast to powerful text-image models, existing audio-image models tend to focus on limited types of natural sounds [6], [7], lacking semantic and emotional context understanding. To overcome this limitation, we treat speech as a special case of sounds differentiated from natural sounds by having content specified in language with an emotional context. We incorporate transcription and speech emotion recognition modules to enhance the understanding of the audio input. The transcribed text is processed the same manner as the text input in the original FRIDA system where it is encoded into the input latent space. The differences between the input features and the features of the planned painting form the loss function for the planner to decide the brush stroke actions. To recognize emotion from the input audio, we use speech emotion recognition (SER) [8] trained with an additional dataset for robustness. We use a pretrained image2emotion model [9] to estimate emotion from the current canvas and then compare that with the emotion recognized from the input audio to form an additional loss term for the planner.

While other sound-guided image manipulation methods are constrained by the dataset used to train their generator model [6], *S-FRIDA* does not rely on a particular pretrained generator and can, therefore, use sound guidance with arbitrary images.

Our surveys show promising results on audio guided paintings, across both natural sounds and emotions as modalities. Furthermore, we qualitatively demonstrate the application of

<sup>1</sup>The Robotics Institute, Carnegie Mellon University  
{vihaanm, pschalde, hyaejino}@andrew.cmu.edu

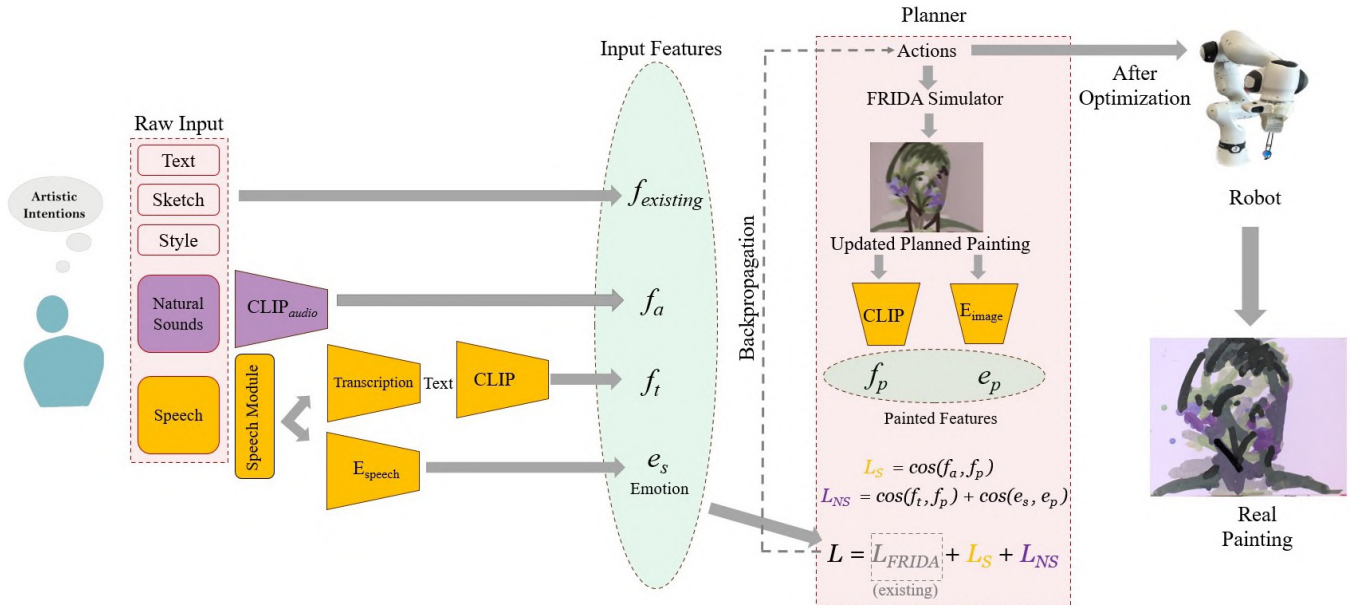


Fig. 2: **S-FRIDA Overview** A human user’s artistic intentions are specified via any combination of natural sounds or speech as well as any of FRIDA’s [1] existing input modalities, e.g., style or sketches. Features are extracted from the audio, or in the case of speech, a transcription and emotions are estimated. Brush stroke actions are rendered into a simulated painting using FRIDA, then features are extracted and compared to the input features to form loss functions. The loss is backpropagated and gradient descent updates the actions to decrease the loss. After optimization, the brush stroke actions are executed by a robotic arm (UFactory XArm and Franka Emika platforms). We observe a high degree of fidelity between the simulated and the real painting drawn by the S-FRIDA framework.

these results through a series of simulated as well as real-world paintings.

Our contributions are as follows:

- 1) We developed a method for sound-guided image manipulation that generalizes across different content types in the painting domain, offering a novel way to integrate auditory cues into visual artwork creation.
- 2) We crafted a speech-to-painting approach that separates speech input into content-interpreted through text-and mood-understood via speech emotion. This method aims to capture both the literal and emotional essence of the spoken input in the artwork.
- 3) To support further research and collaboration, we will make our code publicly available upon the acceptance of this work, hoping to contribute to the broader community’s efforts in this area.

These efforts are intended to provide foundational steps toward leveraging sound and speech as well as emotions as inputs for robotic painting, aiming to capture a richer palette of human expression in art created by robots.

## II. RELATED WORK

### A. Sound-Image Encoding

Encoding sounds and images into the same latent space in a differentiable manner allows for comparison of the two modalities that can be exploited such that one of the modalities can be altered to match the other. [6] train a sound encoder,  $CLIP_{Audio}$ , to encode sound and images into the same latent space.  $CLIP_{Audio}$  was trained on the VGG-Sound Dataset [10] which contains over 200k clips for 309 different classes. The dataset contains 200,000 10-second audio clips

from each class captured from YouTube videos, with no more than two clips per video. The dataset’s sound categories can be broadly divided into: people, animals, music, sports, nature, vehicles, homes, and tools, among others. While  $CLIP_{Audio}$  performs well for simple natural or abstract sounds such as the sound of crying or thunderstorms, it was not designed to understand spoken language and its semantic meanings.

Building on the concept of multimodal learning, Wu et al. [7] introduced Wav2CLIP, a method that distills robust audio representations from the Contrastive Language-Image Pre-training (CLIP) model. By leveraging the visual model of CLIP, Wav2CLIP trains an audio encoder to project audio into a shared embedding space with images and text, facilitating direct comparison and retrieval across modalities.

### B. Emotion-Guided Image Generation

Although CLIP based audio-visual models show good results in image generation, few existing works have studied emotions. [11] used a trained neural network as a feedback function, similar to S-FRIDA, however, the dataset the model was trained on was only 1000 images, none of which were from an artistic domain.

### C. Sound-Guided Visual Manipulation

Many previous works have added sound-guidance to existing image generators to perform sound-guided image or video manipulation. Soundini [12] uses sound to edit videos, and MM-Diffusion [13] generates audio-video simultaneously. The latent space of pre-trained image generators, such as Style-GAN, has been manipulated to understand sound as

input [6], [14]. In other works, music can guide the image generator [15] or perform style transfer [16].

While there has been success in using this methodology for image manipulation, previous work fails to create generalizable output images fitting the sound descriptions because they utilize pre-trained image synthesis models that can only generate images within their training distribution such as churches, faces, or artwork. Therefore, these prior works use sound only to manipulate a particular image that the synthesis model is capable of producing. In this work, our approach creates in a more general content domain and can generate images purely from audio inputs.

### III. BACKGROUND

The proposed approach adds direct speech, sound, and emotional guidance to the FRIDA robotic painting system [1]. Brush strokes form the action space, in FRIDA, which are parameterized by values representing stroke length, bend, and thickness as well as position, orientation, and paint color. Given a set of brush stroke actions and a picture of the canvas, FRIDA’s simulation environment can differentially render the strokes and layer them onto the canvas image forming a simulated painting represented as an RGB image. A loss value can be computed using the simulated painting, and due to the differentiability of FRIDA’s simulation, gradient descent can be used to optimize the brush stroke parameters to minimize the loss.

### IV. APPROACH

Our approach, depicted in Fig. 2, adds sound input to the FRIDA system. Because sounds have vastly different categories, we create approaches for two categories of sound: 1) natural sounds, which are diverse sound samples from various sources, and 2) speech sounds, which are a specific subset of natural sounds made by humans composed of language and tone.

#### A. Natural Sounds Guidance

For natural sound guidance, we use the pretrained  $CLIP_{audio}$  [6] audio-image encoder trained on a wide variety of labelled sounds from YouTube videos. We encode the simulated painting,  $p$ , and input audio,  $a$ , into the same latent space using  $CLIP$  [17] and  $CLIP_{audio}$  respectively. The encodings are compared using cosine distance to form a loss function. In practice, the painting is augmented using various perspective warps and cropping, as is customary in CLIP-guided image synthesis, for robust loss backpropagation.

$$f_a = CLIP_{Audio}(a); \quad f_p = CLIP(p) \quad (1)$$

$$L_{NS}(p, a) = \cos(f_p, f_a) \quad (2)$$

ArtEmis	Ravdess	Crema	TESS	SAV	IEMOCAP
amusement	happy	HAP	happy		hap, exc
anger	angry	ANG	angry	a	fru, ang
awe					
contentment	calm				
disgust	disgust	DIS	disgust	d	dis
excitement	surprise		surprise	su	sur
fear	fear	FEA	fear	f	fea
sadness	sad	SAD	sad	sa	sad
something else	neutral	NEU	neutral	n	neu, xxx, oth

TABLE I: Correspondence between emotion labels in different datasets. From left to right: [4][19][20][21][22][23]

#### B. Speech Guidance

Speech is an special subset of sound that humans use to express their ideas and feelings. We focus on two distinct features of speech, i.e., language and tone to represent the content and its additional emotional context, respectively. Our approach to speech guided painting decouples speech into text and emotions. While  $CLIP_{audio}$  is a promising tool for natural sounds, we find that it is not suitable for comprehending rich semantic meanings of spoken language. To utilize the semantic meaning of the text input, we leverage the text-image  $CLIP$  instead and interpret the audio input separately for the moods.

The input speech is transcribed to text,  $t$ , using the Whisper [18] model. This text can then guide the painting using FRIDA’s existing text-to-painting methodology which encodes both the text and simulated painting into the same latent space using CLIP ( $f_t$  and  $f_p$ , respectively). The cosine distance in features forms a loss function as seen in the first term of Eq. 4.

$$f_t = CLIP(t); \quad f_p = CLIP(p) \quad (3)$$

$$e_p = E_{img}(p); \quad e_s = E_{speech}(a)$$

$$L_S(p, a) = \cos(f_p, f_t) + \cos(e_s, e_p) \quad (4)$$

We adapt an existing framework for Speech Emotion Recognition (SER) trained on four datasets containing speech with labelled emotions [8]. Features such as the Mel Frequency Cepstral Coefficients (MFCC), Mel Spectrogram, and Chromagram are extracted from the input speech waveform audio. A convolutional neural network is trained to predict the emotion. For robustness, we incorporate a fifth dataset for training; The Interactive Emotional Dyadic Motion Capture (IEMOCAP) [23] dataset containing over 12 hours of audio-visual data of people acting along with emotions annotated. Not all of the datasets contained the same set of emotion labels, and some datasets used different terminology for similar emotions. We summarized how we dealt with these emotional label correspondences in Table I.

We guide the painting to have an emotional appearance with the second term of the loss function detailed in Eq. 4. An image-to-emotion prediction model [9] ( $E_{img}$ ) trained on the ArtEmis dataset [4] which contains 80,000 artworks with labeled emotions predicts the emotion probability vector

	Drill	Explosion	Fire	Raining	Thunder	Thunderstorm
Drill	<b>3</b>	1	0	0	0	1
Explosion	0	<b>2</b>	1	0	2	0
Fire	0	0	<b>2</b>	1	2	0
Raining	0	1	0	<b>1</b>	1	2
Thunder	2	0	0	2	<b>0</b>	1
Thunderstorm	0	0	0	0	0	<b>5</b>

TABLE II: The confusion matrix from our survey where participants listened to audio then decided which of the six paintings in Fig. 3 was generated using that audio. Rows display true labels and columns show the predicted labels.

of the simulated painting,  $e_p$ . The emotion from the input speech audio is extracted using our adapted Speech Emotion Recognition model,  $E_{speech}(a)$ . The cosine distance between these features forms the emotional guidance loss.

### C. Multi-Modal Guidance

FRIDA makes painting plans by optimizing the brush stroke parameters to achieve an objective which is made up of the weighted sum of loss functions. Eq. 5 displays FRIDA’s objective function with our two losses  $L_{NS}$  and  $L_S$  included for natural sound and speech guidance, respectively.

$$\tilde{p} = \min_p \left[ w_a L_{\{NS|S\}}(p, a) + \sum_{i=1}^4 (w_i l_i) \right] \quad (5)$$

FRIDA’s base model introduces 4 other loss functions,  $l_i$ , that connect the paintings to modalities text, images, sketches, and styles. The full objective, Eq. 5, is to find the painting plan,  $p$ , that minimizes the weighted sum of losses, given the weights for each loss function,  $w_i$  and  $w_a$ . The optimal painting plan  $\tilde{p}$  is the plan that minimizes the weighted sum of loss functions.

Fig. 8 shows a gradual progression of the influence of emotion guidance as its loss value weight is increased.

## V. RESULTS

We demonstrate how natural sound or speech can be used to produce paintings that match the semantic context of the input sound and that align with the user’s emotional context. We mainly use simulated paintings for evaluation because *S*-FRIDA’s real world paintings take hours to produce, executing hundreds of brush strokes in sequence analogous to human painting. We, however, note that *S*-FRIDA’s real paintings closely match their simulated counterparts as shown in Fig. 2 due to its Real2Sim2Real technique [1]. Real painting results are showcased in Fig. 1 and Fig. 5.

### A. Natural Sounds Guidance

Fig. 3 shows that, using only audio as guidance, our approach can generate paintings that capture various natural sounds. We conducted a survey to quantitatively evaluate the quality of generated images in Fig. 3 in terms of the

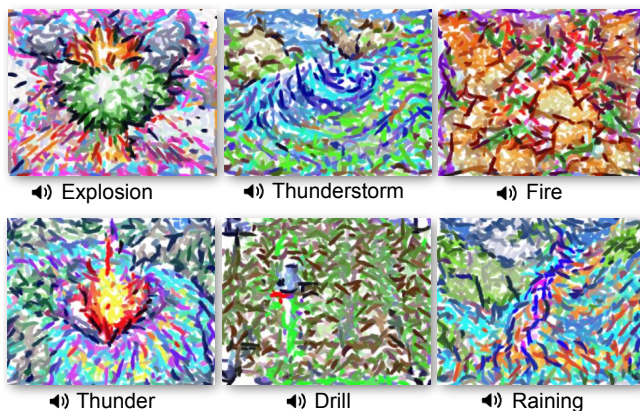


Fig. 3: Paintings generated using natural sounds as input.

human-evaluated semantic similarity between the input audio and the generated painting. In the survey, participants were instructed to listen to an audio sample, then select which of the six images (Fig. 3) were created using that audio as input. A random selection would result in 16.7% accuracy as there were six options; however, participants selected the correct painting 43.3% of the time. We conducted the survey through Amazon Mechanical Turk, recruiting 28 participants. Each audio was evaluated by five different participants. The confusion matrix can be seen in Table II. Many of the sounds were similar and this can be reflected with the frequent confusion of thunder, thunderstorm, and raining sounds.

### B. Emotion Guidance

We first present the quantitative evaluation on emotion guided painting generation, followed by real world paintings as qualitative results.

In this experiment, we examine the quality of generated paintings guided by emotion as a direct input. We use the 9 emotions from the ArtEmis dataset. The first row of Fig. 4 shows *S*-FRIDA’s emotional guidance in isolation where the type of emotion was the only input to produce each painting. The second row of Fig. 4 shows the paintings generated for the same types of emotion jointly with an additional text input, “a house and a tree.”; this result shows the emotional influence on the paintings when their desired contents are described by the users.

Although recognizable content is not generally present in both sets of images, our user survey indicated that the users were able to identify the correct input emotion significantly better when compared to a random guess. In a Mechanical Turk study, we presented participants with one of the eight emotions in Fig. 4 (omitting “Something Else”) in writing along with randomized order of paintings generated using the eight emotions. Participants were asked to identify which painting looks most like the given emotion. The 75 participants were correct 26.5% of the time (random guess would be 12.5%). The most common correctly identified emotion was Awe at 60% correct and the most commonly confused emotions were Amusement incorrectly identified as Excitement 80% of the time.



Fig. 4: First row images: Painting with only emotion guidance. Second row: painting with both emotion and text “A house and a tree.”



Fig. 5: Examples showing how emotion included in the inputs impacts the overall impressions of the paintings. The figure shows the inputs (first row), and the real paintings drawn by S-FRIDA (second row).

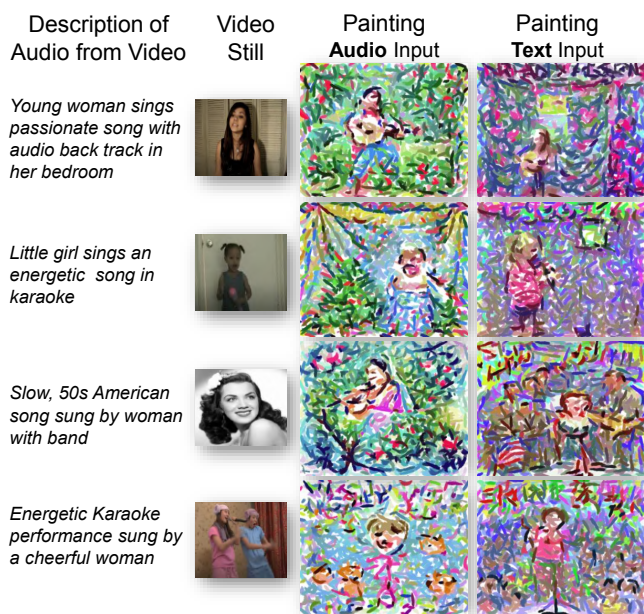


Fig. 6: Painting various examples from the VGG-Sound [10] categorized as “female singing” using audio versus using a description of the audio. A video still is shown here for clarity but is not used in image generation.

Fig. 5 and Fig. 1 (upper) present real world paintings illustrating how emotion—as either a direct input or indirect input as part of speech—affects the impression of the resulting paintings. The first column shows the painting of Andy Warhol using a text input only. The second column shows the painting given the same language description but in a sad toned speech. The third and fourth columns show the paintings given a photograph of Andy Warhol as input together with two specific types of emotion, sad and contentment. All

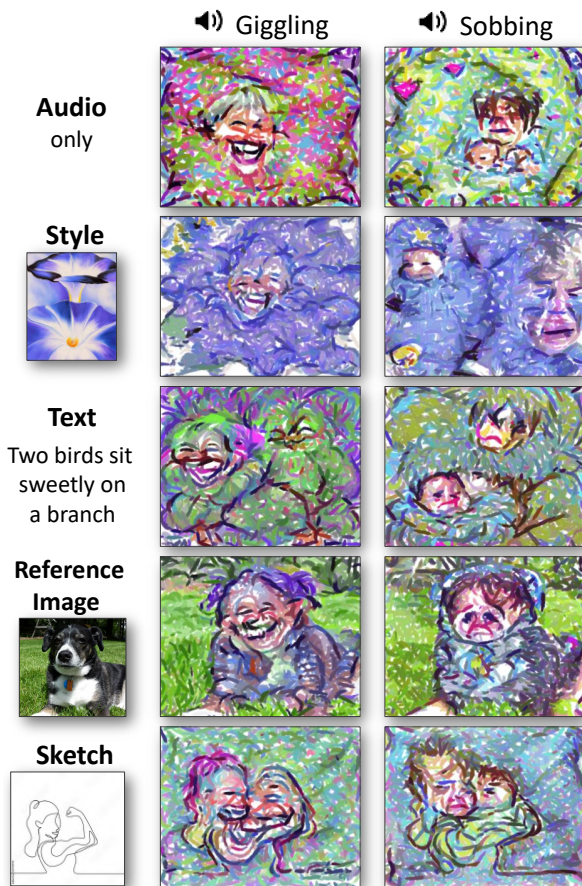


Fig. 7: Combining audio as an input modality with other modalities that FRIDA can handle.

paintings here were painted using acrylic paints in 8" × 10" canvas.

### C. Ineffability of Sound

Sound is nuanced and challenging to describe accurately with language. In this experiment, we qualitatively compare the paintings generated directly from audio with those generated using textual descriptions of the audio as shown in Fig. 6. Specifically, we first sampled multiple examples from the held-out set of the VGG Sound [10] dataset that were all categorized as “female singing.” Given each sample, we generated two paintings: one using the audio as input and the other using the corresponding text description of the audio. While both paintings from audio and text tend to

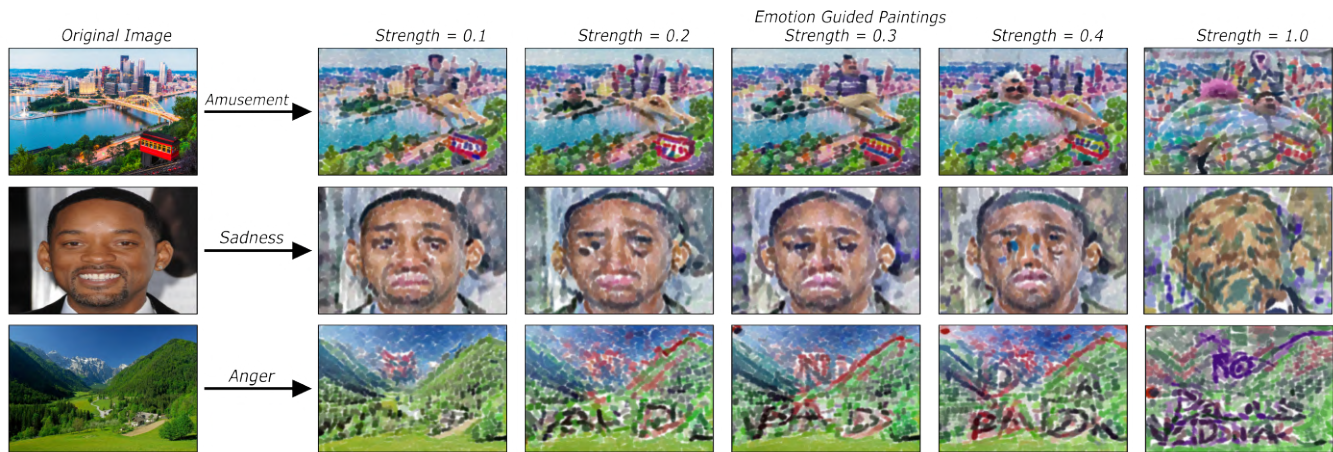


Fig. 8: An image grid depicting the results of emotion guidance with varying guidance strengths. From left to right, each column showcases the progression of the synthesized paintings, with increasing levels of emotion guidance strength, as indicated by the numbers.

align with the content of the video’s sound, they are vastly different from each other. This result hints some evidence that it is challenging to use language to describe the semantic meanings of sounds fully.



Fig. 9: Painting using various Pop songs as input. Genres of the songs from left to right are Disco, Traditional American Folk, and Rap/Hip-Hop

#### D. Additional Results

**Multimodal Guidance:** Paintings were generated with audio inputs paired with FRIDA’s existing modalities in Fig. 7. Loss function weights were adjusted to allow the appearance of both modalities to become prominent. The results appear abstract, although both modalities are strongly present.

**Music Guidance:** The VGG-Sound dataset used to train our sound encoder contains some music in various forms. Music is a modality that is extremely challenging to represent in any other form such as language or images accurately. We experiment with the generalization of the image encoder in Fig. 9. The preliminary results are abstract but make promising resemblance to content and general atmosphere of the given songs.

## VI. DISCUSSION

This work generally contributes toward creative human-robot interaction research. We believe that supporting diverse and intuitive input modalities would endow the users an increased control of a robotic painting system. The users, in turn, would feel more ownership over the artwork that they create collaboratively with the system.

Sound can act as an accessible input modality for people with visual impairments who want to paint. In future work, we hope to engage with different communities of people who

face physical barriers to performing painting to enable more people to express themselves via the visual art of painting.

#### A. Limitations

The generalization of the content that our approach can produce depends on the training data. Our image-to-emotion prediction model is trained on the ArtEmis dataset [4] which pairs paintings with their labelled emotions. Because art data was used, the model is biased towards guiding our images towards this distribution and away from other image distributions such as photographs. While the VGG-Sounds [10] is large with more than 300 categories, it is still not completely general and we observe its limitations in the examples such as music guided paintings.

As in other generative art tasks, evaluating the quality of produced outputs remains challenging.

S-FRIDA’s paintings appear abstract and noisy compared to recent, pixel-generating models, e.g., Stable Diffusion [24]. This is partially due to using CLIP and other neural networks that were trained for image categorization rather than generation, as seen with other methods that use CLIP for image generation [25], [26]. Another source of the abstract appearance is from using simple, brush strokes as image primitives rather than pixels. FRIDA plans with a single brush that is not capable of producing very fine details.

## VII. CONCLUSIONS

Our approach, Robot Synesthesia, adds sound inputs to the FRIDA [1] robotic painting platform. When compared to existing sound-guided image synthesis approaches, ours is more general as it does not rely on constrained pre-trained image generators, e.g., Style-GAN. In addition, we treat natural sounds and speech separately to unlock the important content and emotion held in spoken language. While many existing work focuses on robots following commands, Robot Synesthesia is a rare attempt at a robotic system which can hear a human user, understand their emotions, and assist them to express their ideas in visual art. By supporting audio-based interaction, our work contributes to make a robotic painting system accessible to broader user groups.

## REFERENCES

- [1] P. Schaldenbrand, J. McCann, and J. Oh, "Frida: A collaborative robot painter with a differentiable, real2sim2real planning environment," in *International Conference on Robotics and Automation (ICRA)*, 2023.
- [2] P. Schaldenbrand, G. Parmar, J.-Y. Zhu, J. McCann, and J. Oh, "Cofrida: Self-supervised fine-tuning for human-robot co-painting," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2024.
- [3] N. Kaya and H. H. Epps, "Relationship between color and emotion: A study of college students," *College student journal*, vol. 38, no. 3, pp. 396–405, 2004.
- [4] P. Achlioptas, M. Ovsjanikov, K. Haydarov, M. Elhoseiny, and L. J. Guibas, "Artemis: Affective language for visual art," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11569–11579, 2021.
- [5] F. Masina, V. Orso, P. Pluchino, G. Dainese, S. Volpato, C. Nelini, D. Mapelli, A. Spagnolli, and L. Gamberini, "Investigating the accessibility of voice assistants with impaired users: Mixed methods study," *J Med Internet Res*, vol. 22, p. e18431, Sep 2020.
- [6] S. H. Lee, W. Roh, W. Byeon, S. H. Yoon, C. Y. Kim, J. Kim, and S. Kim, "Sound-guided semantic image manipulation," 2021.
- [7] H. Wu, P. Seetharaman, K. Kumar, and J. P. Bello, "Wav2clip: Learning robust audio representations from CLIP," *CoRR*, vol. abs/2110.11499, 2021.
- [8] S. Burnwal, "Speech emotion recognition," 2020.
- [9] R. Szymczak, "Predicting emotional response to visual stimuli, a machine learning approach," tech. rep., 2022.
- [10] H. Chen, W. Xie, A. Vedaldi, and A. Zisserman, "Vggsound: A large-scale audio-visual dataset," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 721–725, IEEE, 2020.
- [11] F. Ibarrola, R. Lulham, and K. Grace, "Affect-conditioned image generation," *arXiv preprint arXiv:2302.09742*, 2023.
- [12] S. H. Lee, S. Kim, I. Yoo, F. Yang, D. Cho, Y. Kim, H. Chang, J. Kim, and S. Kim, "Soundini: Sound-guided diffusion for natural video editing," 2023.
- [13] L. Ruan, Y. Ma, H. Yang, H. He, B. Liu, J. Fu, N. J. Yuan, Q. Jin, and B. Guo, "Mm-diffusion: Learning multi-modal diffusion models for joint audio and video generation," 2023.
- [14] O. Patashnik, Z. Wu, E. Shechtman, D. Cohen-Or, and D. Lischinski, "Styleclip: Text-driven manipulation of stylegan imagery," 2021.
- [15] D. Jeong, S. Doh, and T. Kwon, "Träumerai: Dreaming music with stylegan," 2021.
- [16] C.-C. Lee, W.-Y. Lin, Y.-T. Shih, P.-Y. P. Kuo, and L. Su, "Crossing you in style: Cross-modal style transfer from music to visual arts," in *Proceedings of the 28th ACM International Conference on Multimedia*, MM '20, (New York, NY, USA), p. 3219–3227, Association for Computing Machinery, 2020.
- [17] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, "Learning transferable visual models from natural language supervision," 2021.
- [18] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," *arXiv preprint arXiv:2212.04356*, 2022.
- [19] S. R. Livingstone and F. A. Russo, "The ryerson audio-visual database of emotional speech and song (ravdess): A dynamic, multimodal set of facial and vocal expressions in north american english," *PLoS one*, vol. 13, no. 5, p. e0196391, 2018.
- [20] H. Cao, D. G. Cooper, M. K. Keutmann, R. C. Gur, A. Nenkova, and R. Verma, "Crema-d: Crowd-sourced emotional multimodal actors dataset," *IEEE transactions on affective computing*, vol. 5, no. 4, pp. 377–390, 2014.
- [21] K. Dupuis and M. K. Pichora-Fuller, "Toronto emotional speech set (tess)-younger talker.happy," 2010.
- [22] P. Jackson and S. Haq, "Surrey audio-visual expressed emotion (savee) database," *University of Surrey: Guildford, UK*, 2014.
- [23] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "Iemocap: Interactive emotional dyadic motion capture database," *Language resources and evaluation*, vol. 42, pp. 335–359, 2008.
- [24] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.
- [25] K. Frans, L. Soros, and O. Witkowski, "Clipdraw: Exploring text-to-drawing synthesis through language-image encoders," *arXiv preprint arXiv:2106.14843*, 2021.
- [26] P. Schaldenbrand, Z. Liu, and J. Oh, "Styleclipdraw: Coupling content and style in text-to-drawing translation," in *Proceedings of the International Joint Conference on Artificial Intelligence*, 2022.