



Exploring the Use of Generative AI in the Search for Extraterrestrial Intelligence (SETI)

John Hoang ^{*(1)}, Zihe Zheng⁽²⁾, Aiden Zelakiewicz⁽³⁾, Peter Xiangyuan Ma⁽⁴⁾, and Bryan Brzycki⁽⁵⁾

(1) SCIPP, California, USA; e-mail: jokhoang@ucsc.edu

(2) Yale University, Connecticut, USA; e-mail: zihe.zheng@yale.edu

(3) The Ohio State University, Ohio, USA; e-mail: zelakiewicz.1@osu.edu

(4) Department of Mathematics, University of Toronto; e-mail: peterxy.ma@mail.utoronto.ca

(5) Department of Astronomy, University of California Berkeley; e-mail: bbrzycki@berkeley.edu

Abstract

The search for extraterrestrial intelligence (SETI) is a field that has long been within the domain of traditional signal processing techniques. However, with the advent of powerful generative AI models, such as GPT-3, we are now able to explore new ways of analyzing SETI data and potentially uncover previously hidden signals. In this work, we present a novel approach for using generative AI to analyze SETI data, with focus on data processing and machine learning techniques. Our proposed method uses a combination of deep learning and generative models to analyze radio telescope data, with the goal of identifying potential signals from extraterrestrial civilizations. We also discuss the challenges and limitations of using generative AI in SETI, as well as potential future directions for this research. Our findings suggest that generative AI has the potential to significantly improve the efficiency and effectiveness of the search for extraterrestrial intelligence, and we encourage further exploration of this approach in the SETI community. (**Disclosure:** For the purpose of demonstration, the abstract and title were generated by ChatGPT and slightly modified by the lead author.)

1 Introduction

The Breakthrough Listen project has been searching for technosignatures in our universe using powerful radio telescopes around the world, including the Green Bank Telescope, Parkes Telescope, Allen Telescope Array, and MeerKAT. The most popular technique in radio SETI involves looking for narrow-band signals in the time-frequency spectrogram data (often referred to as dynamic spectra or waterfall plots in existing literature). In recent years, machine learning (ML) algorithms have been employed to classify image-like spectrograms [1], and `setigen`, an open-source Python library, was created to synthesize mock labelled radio SETI training data set [2]. Since `setigen` is meant to be a general-purpose heuristic framework, there are rooms for improvement if one wishes to improve the speed of the algorithm in specialized cases. The method used in this work is Generative Adversarial Network (GAN) [3].

A GAN consists of two competing deep neural networks: the generative network and the discriminator network. The generator generates images from random noise, and the generated images are fed into the discriminator along with real images. The discriminator then classifies the images as either real or fake and both the generator and the discriminator models are updated according to the result. At equilibrium, the generator will generate images that look similar to the training data and as a consequence, the discriminator will no longer be able to distinguish between generated and real images. The output of the generator will determine whether an image is similar to the training data or not, which can be useful in determining whether the image is normal or anomalous. Therefore, GAN can also be used to detect outliers. In astronomy, GANs have been used to simulate galaxy images [4], gamma-ray Cherenkov air-shower signals [5], detecting outliers [6], to name but a few. GANs and its family of other generative frameworks such as AutoEncoder are parts of what is colloquially known as DeepFake, a portmanteau of Deep learning and Fake.

2 Experiments

We employ the `setigen` software to generate waterfall plots used for training. Each generated set includes 15000 waterfall plots, each having 128 pixels in time and 128 pixels in frequency. Background noise is randomly chosen from either Gaussian noise or Chi-squared noise, with mean equals to 10 and standard deviation (only for Gaussian noise) equals to 1. Injected narrow-band signals are specified by the starting point, drift rate and width in each data set. The width of the signals all range from 50 to 60 pixels. Random walk signals are injected in some of the data sets to simulate radio-frequency interference (RFI), which is frequently seen in real observations. These 4 sets used in the experiments are shown in Fig.1.

2.1 Simple drifting signals

We first test GAN's ability to generate waterfall plots containing narrow-band signals with a single drift rate. As seen from Fig.2, GAN successfully reproduces this data set after training.

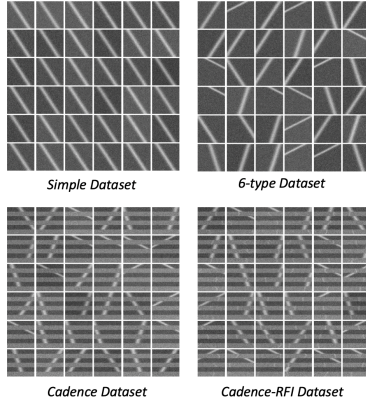


Figure 1. 4 narrow-band data sets used for training in this work with increasing complexity. Top row: narrow-band signals with a constant drift rate, narrow-band signals with 6 drift rates. Bottom row: ON-OFF cadence observations with narrow-band ETI-like signals in only ON observation, ON-OFF cadence observations with narrow-band ETI-like signal in only ON observation and random narrow-band RFI-like in OFF observation.




Figure 2. Vanilla GAN is able to reproduce the training data from `setigen` (left): right plot contains GAN-generated waterfall plots with random background noise level

2.2 Simple drifting signals with varying drift rates

We then expand the variety of data by changing the drift rate and starting points of signals in the waterfall plots, which result in the 6-type dataset. We train the same vanilla GAN network architecture on the 6-type data set shown in Fig.1 for 1000 epochs. However, the network fails to produce realistic waterfall plots.

Conditional GAN [7] is then selected as our next candidate architecture among several variations of GAN since it allows the network to generate images conditioned on the class of data. Conditional GAN achieved this idea by adding labels for data according to their drift rates, which are fed in the network along with the plot themselves. We labeled the 6-type data set with 0, 1, 2, 3, 4, 5 according to 6 different drift rates, and trained the conditional GAN on the data set for 1000 epochs. The generator from the conditional GAN is able to generate waterfall plots with a singular signal with various starting points in each class.

The results in comparison with traditional GAN are shown in Fig.3.




Figure 3. Vanilla GAN struggles to learn the distinct features among different drift rates. However, conditional GAN solves the issue by artificially condition training data on the slope label.

2.3 Cadence signals

Due to the prevalence of radio frequency interference, a typical radio SETI observation consists of pointing the telescope alternatively to ON-OFF sky locations, and a potential narrow-band ETI signal candidate should only show up in the ON observations. To simulate this, three horizontal noisy band of size 21×128 , 21×128 and 23×128 pixels are added to each waterfall plots to simulate the ON-OFF cadence in real observations. Successful results are shown in Fig.4




Figure 4. Conditional GAN is able to learn the more complex ON-OFF signal patterns mimicking typical SETI observations.

2.4 Cadence signals with RFI

We finally test our generator on its resistance to RFI in the dynamic spectra. We trained our conditional GAN model on the Cadence-RFI dataset, which contains injected RFI noise. The result is shown in Fig.5. The generator ignores the presence of RFI and does not produce them in the OFF observations. In addition, the generator also avoids a common failure mode called mode-collapse whereby GAN only learns to output a particularly plausible output and nothing else. Fig.5 shows a degree of variability within a single class: although all the waterfall plots in the same column have the same drift rate (label), they vary with intensity and location within the frame.




Figure 5. Conditional GAN is able to ignore RFI in the training data. For clarity, in the right plot, waterfall plots in each column have the same label (drift rate).

2.5 Generator’s scores

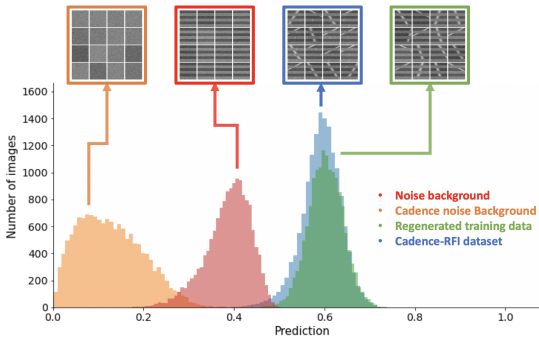


Figure 6. Distribution of scores for different data sets by the discriminator.

Because the discriminator learns to classify realistic waterfall plots and unrealistic waterfall plots during training, the trained discriminator can be used subsequently as a classifier. We train our conditional GAN on the cadence data set in Fig.1. To test the trained discriminator, we used the 4 data sets shown in Fig.6, each containing 15000 waterfall plots: only noise background, only cadence background, GAN-produced data, and the cadence-RFI data. It is clear that waterfall plots that are less similar to training data receive lower scores from the discriminator than those are similar to the training data set. The blue and green distribution have similar means, confirming that both generator and discriminator disregard RFI in the training data.

2.6 Timing

To illustrate the advantage of using pre-train model, we show in Tab.1 the time it took our GAN to generate and write to storage 10^2 , 10^3 , 10^4 , 10^5 waterfall plots. Comparing with the clearly linear time of `setigen`, GAN’s speed is superior when a large data set is involved, and we are only limited by the availability of GPU RAM and storage I/O.

3 Possible extensions

It is possible to further modify GANs into a more advanced bi-directional architecture so that we can reverse-engineer

Table 1. CPU/GPU time to generate (G) and write (W) to storage as a function of number of waterfall plots, using both heuristic `setigen` with CPU and generative GAN with GPU. GAN excels when a large number of waterfall plots are involved, improving the speed by more than 2 orders of magnitude.

#	File Size	SETIGEN		GAN	
		G (s)	G&W (s)	G (s)	G&W (s)
10^2	12.5 MB	0.578	0.613	3.335	3.378
10^3	125.3 MB	5.09	6.287	3.784	4.181
10^4	1.2 GB	51.279	55.723	3.807	7.724
10^5	12.2 GB	515.802	563.833	4.358	95.597

images into meaningful latent vectors. The proposed Bidirectional Conditional GAN (BiCoGAN) [8] includes an additional encoder network that is trained simultaneously with the generator and the discriminator, and can provide inverse mappings of data samples to both intrinsic and extrinsic vectors. The benefit of BiCoGAN is two-fold: the reverse-engineered latent vectors of hitherto unseen signals can be used to compare with those of known signals to identify anomalies, and anomalous latent vectors can be slightly perturbed and subsequently fed into the generator to synthesize new anomalous signals useful for further training. This remains a work in progress, and an example of hand-written digits generated by BiCoGAN is shown in Fig.7.




Figure 7. MNIST images synthesized by a BiCoGAN’s generator with random latent vectors (top row), comparing with original training data (middle row) and with reconstructed latent vectors from the encoder network (bottom row).

4 Limitations

For the purpose of demonstrating the limitations of Generative AI, let us examine this paper’s title and abstract, which were largely produced by ChatGPT - Chat Generative Pre-trained Transformer. While reads convincingly at first, the generated abstract certainly contains errors. Notably, the relationship between natural language processing and radio data analysis makes little sense, if at all. Nonetheless, upon a handful of minor edits by an expert, the hybrid version of the abstract sounds even more convincing (see Fig.8). In a similar vein, radio spectra produced by a GAN are realistic overall, but they require more careful benchmark when high precision is required or when statistics are low.

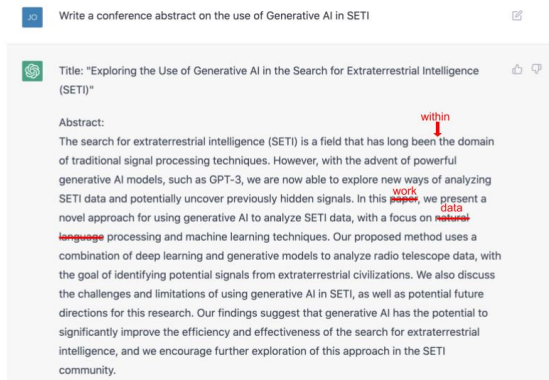


Figure 8. Response by ChatGPT from a prompt by the author, and with subsequent edits to either improve language usage or correct factual knowledge.

Fundamentally, by the virtue of the Universal Approximation Theorems, neural networks are functional approximator (often polynomial approximation if simple activation functions are used) of a multidimensional distribution to arbitrary precision. In essence, it is analogous to the Weierstrass Approximation Theorem, which is as follows: Suppose $f(x)$ is a continuous real-valued function defined on the real interval $[a,b]$, then:

$$\forall \varepsilon > 0, \exists p(x) \text{ s.t. } \forall x \in [a,b], |f(x) - p(x)| < \varepsilon \quad (1)$$

In the case of ChatGPT or any Generative AI, the interval $[a,b]$ is the domain in which the training data $f(x)$ is contained. Since the domain of the function is a specific *closed interval* instead of the entire real number line, the generator can approximate or interpolate very well the distribution *within the domain* using a polynomial $p(x)$. In other words, the output is a sufficiently good polynomial approximation of the training data contained therein. In the case of this work's abstract, since ChatGPT most likely obtained training data from publications, it outputs "paper" instead of the more naturally sounded "work" in the context of an abstract for a conference presentation. However, *outside the domain*, the *extrapolated* output is sometimes only tangential, and requires further scrutiny. The ideal training data set should include both theory-based data as well as real-life examples, although this is hardly achievable in many cases whereby new phenomenon are being investigated.

5 Conclusions

We have explored the potential of deep generative networks for generating dynamic spectra and their ability to detect outliers. The main advantage of such generative AI is that they can encode the complex representations of the training data and rapidly decode them at a rapid speed by leveraging GPU usage. The work is relevant not only to radio SETI, and can be readily extend to SETI search in other parts of the electromagnetic spectrum. Furthermore, since the AI models used in this work is still far from the industrial grade models such as one used by ChatGPT,

there is much for further experimentation with more sophisticated AI architecture. However, as in the case with other applications of Generative AI, we caution against the over-reliance on their outputs, as sometimes they are only tangential to the reality at best, and nonsensical at worst. The code used in this work is publicly available at <https://github.com/zzheng18/SETIGAN>.

Acknowledgements

We thank the Breakthrough Prize Foundation and the University of California, Berkeley, for their support. Special thanks to Andrew Siemion, Steve Croft, and Yuhong Chen for several fruitful discussions. This work is funded by the Breakthrough Listen Initiative and the National Science Foundation.

References

- [1] P. X. Ma et al., "A deep-learning search for technosignatures from 820 nearby stars," *Nature Astronomy*, pp. 1–11, Jan. 2023, doi: <https://doi.org/10.1038/s41550-022-01872-z>.
- [2] B. Brzycki et al., "Setigen: Simulating Radio Technosignatures for SETI," *The Astronomical Journal*, vol. 163, no. 5, p. 222, May 2022, doi: [10.3847/1538-3881/ac5e3d](https://doi.org/10.3847/1538-3881/ac5e3d).
- [3] I. J. Goodfellow et al., "Generative Adversarial Networks," *arXiv.org*, 2014. <https://arxiv.org/abs/1406.2661>
- [4] M. Dia, E. Savary, M. Melchior, and F. Courbin, "Galaxy Image Simulation Using Progressive GANs," *arXiv:1909.12160 [astro-ph, stat]*, Sep. 2019, Accessed: Jan. 25, 2023. [Online]. Available: <https://arxiv.org/abs/1909.12160>
- [5] J. Dubenskaya, A. Kryukov, and A. Demichev, "PoS(ICRC2021)874 Fast Simulation of Gamma/Proton Event Images for the TAIGA-IACT Experiment using Generative Adversarial Networks PoS(ICRC2021), Accessed: Jan. 25, 2023. [Online]. Available: <https://pos.sissa.it/395/874/pdf>
- [6] B. Margalef-Bentabol et al., "Detecting outliers in astronomical images with deep generative networks," *Monthly Notices of the Royal Astronomical Society*, vol. 496, no. 2, pp. 2346–2361, Jun. 2020, doi: [10.1093/mnras/staa1647](https://doi.org/10.1093/mnras/staa1647)
- [7] M. Mirza and S. Osindero, "Conditional Generative Adversarial Nets," *arXiv.org*, 2014. <https://arxiv.org/abs/1411.1784>
- [8] MA. Jaiswal, W. AbdAlmageed, Y. Wu, and P. Natarajan, "Bidirectional Conditional Generative Adversarial Networks," *arXiv:1711.07461 [cs, stat]*, Nov. 2018, Accessed: Jan. 25, 2023. [Online]. Available: <https://arxiv.org/abs/1711.07461>