



SnakeChat: a conversational-AI based app for snake classification

Jorge Guerra Pires

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.

Abstract

Conversational Artificial Intelligence (CAI) was the latest big step towards artificial general intelligence (AGI): models of AI that are generic, that can solve problems without being restricted to a small set of problems they were designed for. OpenAI with chatGPT gave a big step towards this direction. On the other hand, being able to classify snakes from images can help the general public, and even experts to classify snakes (e.g., finding the right snake serum). There are more than 2.000 different species of snakes worldwide. On this paper, we explore the latest releases of openAI APIs. We integrate with a previous work we did, using transfer learning to classify snakes. We have used their GPT Vision API for describing uploaded images textually, their latest chatGPT API for deciding which model to use for classifying an image and also to generate a human-friendly final textual response. Additionally, we also tested their GPT 4 API instead of the chatGPT API as a comparison. We found that the integration was a success. By giving our models to openAI API, it decides which model to call as so it could give an educated guess to the user based on an uploaded image from the snake. The best result was when we used the GPT 4 API, something we already expected. We hope that this prototype can be scaled up with more species. We did a test of concept, and showed that the basics needed to build such a system for snake classification is possible with current technologies. We hope to contribute to sparking research that brings together computer scientists and biologists, generally called bioinformatics.

Jorge Guerra Pires*

Founder at IdeaCoding Lab / JovemPesquisador.com, Brazil

*Email address: jorgeguerrapires@yahoo.com.br

Keywords: biology; bioinformatics; conversational artificial intelligence; Snakes; openAI; chatGPT; chatbots.

1. Introduction

When chatGPT was first released, we could not believe what we were seeing: an artificial intelligence with high -language capabilities; they are called *conversational artificial intelligence*, conversational AI for short. So far, we were used to models that talked in experts' language, we needed those experts to translate to the public their outputs. They were very complex and specialized.

Conversational artificial intelligence (CAI) is a branch of computer science that aims to create systems that can interact with humans using natural language.

CAI systems can be used for various purposes, such as customer service, education, entertainment, health care, and more. CAI systems can be classified into two types, mainly: task-oriented and chat-oriented.

Task-oriented CAI systems are designed to help users accomplish specific goals, such as booking a flight, ordering a pizza, or checking the weather. Chat-oriented CAI systems are designed to engage users in open-ended conversations on various topics, such as sports, movies, or hobbies. Chat-oriented CAI systems are more challenging to develop, as they require a deeper understanding of natural language, human emotions, and common sense.

chatGPT, as it was firstly released, is a chat-oriented CAI, the first to gain popularity; in fact, we could manually make it also a task-oriented since it is a generic model. Recently, openAI launched versions of chatGPT that are also task-oriented CAI, therefore, the line is becoming fuzzy, hard to separate where starts one type and where begins the other. Our system, thanks to this hybrid-behaviour of openAI API, is a hybrid CAI (Task-oriented + Chat-oriented). In fact, before chatGPT, chabots would require readings books and books, and coding a lot. Now, one can create one, state-of-the-art in hours or less, see for learning to build one using Wix and chatGPT API Pires (2023b).

We are going to use openAI APIs. An API, or application programming interface, is a set of rules and specifications that define how different software components can communicate and interact with each other. APIs are essential for building complex and scalable applications, as they allow developers to reuse existing code and functionality, rather than reinventing the wheel everytime. APIs can also enable integration between different systems and platforms, such as web services, databases, cloud computing, mobile devices, etc. APIs can be classified into different types, such as RESTful, SOAP, GraphQL, etc., depending on the format and protocol they use to exchange data; the basic protocols we are using is JSON format as output from our functions and APIs we use. APIs can also have different levels of security, documentation, and testing, depending on the needs and requirements of the users and developers.

Public API in artificial intelligence was an important paradigm shift. The development and adoption of API in artificial intelligence has been an important paradigm shift in the field of computing and software engineering. API in artificial intelligence enables developers and users to access and leverage the power of advanced machine learning models and algorithms without having to build them from scratch or understand their inner workings.

OpenAI APIs are quite simple: text in text out; they have some basic parametrizations, but are very simple compared to classical expert-driven APIs. This reduces the cost, time and complexity of developing and deploying AI solutions, and democratizes the access and use of AI across various domains and industries. API in artificial intelligence also fosters innovation and collaboration, as different AI services can be combined and integrated to create new functionalities and applications. Some examples of popular AI APIs are Google Cloud Vision API, IBM Watson API, Amazon Rekognition API, Microsoft Cognitive Services API, etc. and more recently, openAI APIs. We are actually gathering two artificial intelligence APIs herein: TensorFlow.js from Google, behind SnakeFace, and openAI APIs.

We are going to focus on a task-oriented CAI system, once the main goal of our system is classifying snakes, by using another system called SnakeFace Pires and Dias Braga (2023), that we designed previously using transfer learning. Of course, future versions can be more chat-oriented by providing means for the user to learn more, to ask for more information, to have a "snake-friendly conversation". openAI CAI can be a very rich and versatile source for learning more about any topic; their latest model *gpt-4-1106-preview* was trained on information until Apr 2023.

Thanks to how the openAI APIs are built, it is easy to change the focus of the system. For creating a task-oriented CAI, we are going to explore a new feature of their API called *function calling*: you give the algorithm a set of functions to be used, and the system will decide which one to use. On this approach, as one example, a system could decide to call an external API on traffic to decide whether it should recommend a specific route. In our

case, we are going to call external models to confirm a textual description from an image from a snake. This description can be done by the user, or a second model from openAI will create the image description using GPT 4, the sequential model of chatGPT (chatGPT is powered by GPT 3.5).

Until now, AI models were designed for programmers who knew how to read that 0.8 means with a probability of 80%, and that the model could make mistakes. Artificial intelligence was for experts. chatGPT made artificial intelligence to gain the common world. Before chatGPT, language models were quite limited, such as content moderation; and large language models were limited to their own companies that developed them. Now, they could create complex and long texts. chatGPT is what they call a Large Language Model (LLMs), their first GPT version included more than 80 GB of texts, which increased in the later versions Kulbik and Saboo (2022). Certainly, people were intrigued by with how chatGPT worked Wolfram (2023). One widely applications is for creating chatbots Pires (2023b).

Before chatGPT, Google already had advanced AI APIs, made available for programmers, some paid, some free of charge for usage; SnakeFace was built using Teachable Machine, an AI platform using transfer learning made available by Google free of charge. What was done different by openAI is making available an advanced AI model with low cost, a large language model (LLM). OpenAI makes available the pieces one needs to develop their own advanced apps based on AI. They create their own apps, but also making available the tools they used. This is a paradigm shift: creating APIs instead of the finished software; giving the bricks, instead of the pre-built house, build your own house, you own style. This new paradigm creates a new industry: people can create their own advanced AI apps, without actually creating an AI model. This is certainly a new way of selling artificial intelligence. It even created what some called "the industry of prompts": people can make money by creating their own textual commands to chatGPT, and embed that on a production software Kulbik and Saboo (2022).

We suspect, as we used their models, that their biggest revolution was chatGPT; the sequentials such as GPT 4, that we are going to explore herein, are applications of GPT 3.5, the algorithm behind chatGPT. For instance, their vision model does not seem to be better than other vision models already available. As we are going to see, it seems to make the same mistakes as other computer vision models we used do. We suspect they used the same trick we did: feed a LLM with information from another AI model (e.g., MobileNet). Externally, may seem like something different, but the trick is one single big step can serve as leverage to new big steps. Thus, even though some may suggest that GPT 4 sees different from other computer vision models, we suspect they are using classical model for vision, and connecting to their language model, which was indeed a revolution. Their revolution was on language models, not on computer vision, or other tools they are offering, such as data science (e.g., code interpreter).

One interesting observation we are going to see: they do not seem to be good at "reasoning", handling conflicting information and arrive to the right conclusion. Which can be seen a *deductive reasoning* capability. As one example, in one response from openAI API, all the correct information was on the response, but the final response was wrong. That was seen also when the image description was not describing a snake species that the model guessed, that would be impossible if one compares the description with the snake species coloration and patterns. This would be an ability to handle conflicting information, and arrive to the right conclusion. Something we human can do easily.

The same behavior was seen on BARD and Bing Chat. It could be explained if we assume those AIs are trained to mine/generate information/text, rather than reasoning on information/text. As we are going to see: they do not seem to be good at conflicting information from different sources. They may even create a base reasoning for a wrong prediction. Those are certainly undesirable behaviors we need to pay attention to.

Large language model is a type of artificial intelligence system that can generate natural language texts based on a given input, such as a word, a phrase, or a prompt; they can also execute text-related tasks, such as answering questions or summarizing information. Large language models are trained on massive amounts of text data, such as books, articles, websites, and social media posts, to learn the patterns and structures of natural language. Large language models can perform various tasks, such as answering questions, summarizing texts, writing essays, creating captions, and more. However, large language models also have limitations and challenges, such as ethical, social, and environmental implications.

One challenge chatGPT faced was hallucination Emsley (2023): it would make up the answer if it could not answer. One solution that partially solved the problem is allowing chatGPT to call external sources of information. In some cases, as Bing, it is the browser, for others, it was powerful software Wolfram (2023). We follow a similar path: we give chatGPT the ability to call external functions we designed for snake identification. See that both the SnakeFace and SnakeChat are prototypes, for production usage, they must be scaled up, adding a considerable amount of snakes. They are at the current state tests of concepts.

On this article, we explore the latest APIs from openAI. We have tested how it would support on a vision we had on our previous paper Pires and Dias Braga (2023). We have created a set of small models for snake classification, and we envisioned that chatGPT would be useful for actually choosing the model to use. Since then, openAI launched new models. One of these models, called function calling, is what we needed. Another release we are going to explore is the API version of GPT-4 vision capability. We are going to use it to describe an entered image, in case the user does not provide a description for the image. The description is used to decide which model to use (the function calling capability we mentioned).

1.1. Conversational Artificial Intelligence in biology

We have done a research using Semantic Scholar and Bing powered by chatGPT, and we were unable to find papers on chatbots in biology. Thus, to the best of our knowledge, we are the first researchers to apply openAI APIs to biology on the context of scientific research. However, we will still keep searching, and possibly publish a review paper if we find a considerable amount in time. The tendency is that they will be more and more common, as we are doing. Once the cost for implementing those CAI dropped, also, the entering barrier will also drop.

Building a CAI by itself would require a big team, but thanks to openAI, we can just use it, as we are doing herein; it is hard to imagine academic groups having access to such amounts of investments to build such a CAI. That could explain why we were unable to find papers on this topic: it would require a huge team to create the CAI, and then apply it to biological cases; the investment barrier was high. This is no longer the case. This is just one book on CAI before openAI APIs Freed (2021).

Which means, we are innovating. We also tried to find CAI in bioinformatics, and the results are focused on medicine. We believe that we may spark researches on using those bots in biology, which could help on using those models in biology. Most models can be intimidating, and as we are going to see, we replace all the user interaction with SnakeFace by a chatbot, which we called SnakeChat.

Chatbots can be the UI/UX for the future of biology and medicine. We found on an informal survey that parametrization of models is a big barrier for medical doctors and biologists when using models Pires (2021).

1.2. A brief look at the literature on snake classification

Transfer learning applied to snake classifications is an important and promising technique in the field of machine learning Kalinathan et al. (2021); Mirunalini et al. (2022); Yang and Sinnott (2021); Proggia et al. (2021); Lakshmi et al. (2021); Krishnan (2020). Several papers have explored this area, highlighting its potential for snake species identification and classification. SnakeFace Pires and Dias Braga (2023), a transfer learning based app, demonstrates the ease of building machine learning models without coding expertise and its adaptability to different biological image identification cases. Another paper proposes a deep learning model that uses transfer learning to build a snake species classifier, combining snake photographic images with geographic location information Desingu et al. (2021).

The problem of snake classification is the problem of assigning classes to snakes based on images. Those assignments can be such as their species. This is a problem of *computer vision*. This problem may have several applications, such as educational purpose, finding the proper venom serum. It is crucial to get the snake correctly, at least, with a certain level of accuracy.

Nonetheless, this problem may impose several challenges. We may have images that are not clear, or images that are not enough for forming a pattern on the training. In other cases, it is possible that the necessary information is not on images. See Pires and Dias Braga (2023) for more discussions about SnakeFace, one of our base models.

One challenge we are going to face is that as the model becomes generic with more species on a single model, also, their precision on species starts to go down. However, as you have models with high accuracy, they become susceptible to bad usage (e.g., entering an image of a snake not present at the model). And this is a trade-off hard to balance, as we are going to see,

Our model stands out on two main points: it was built in JavaScript and it was built to be used. The models we found on the literature are designed in Python Rajabizadeh and Rezghi (2021); Kalinathan et al. (2021); Yang and Sinnott (2021); Desingu et al. (2021); Durso AMand de Castaneda R (2021); Felipe Guimaraes Dos Santos (2021), and they are not available for being easily used. This dominance of Python in machine learning is well-

known, but challenged recently by JavaScript. There are several advantages of using JavaScript, the main one: web applications are built in JavaScript. Thus, one language for all the process: from machine learning to the app development. It may reduce for instance costs since you need just one language, and sometimes, even no server (e.g., Node.js or PHP).

If those models in Python were available as APIs, or any other easy form to be used (e.g., NPM packages), we would have been happy in comparing our findings with other models for snake classification. For so, we would need to implement one by one. It does not seem practical. Nonetheless, SnakeFace showed to be competitive with those models Pires and Dias Braga (2023). None of those models, to the best of our knowledge, explore chatbots. Chatbots can make it easier to use models, since all they require is a human-like command, and the artificial intelligence will handle the rest.

The reader can find a more detailed discussion about the literature on snake classification on the paper about SnakeFace Pires and Dias Braga (2023).

It is important to mention that Instaturalist, from where we sampled most of our images of snakes, has their own open source app project, which is a general app. Their project is a community project. I have used their app for suggesting snakes: it tends to be good at associating similar snakes. Nonetheless, it is not good at actually make a final prediction. Initial prediction are later double-checked by humans, their community. The final prediction is a community consensus. Thus, it is an annotation process, more precisely, a crowd annotation. Their platform is perfect for us since the images are real, they are not professional images. The images is the type of images users may take when using our app, amateur images.

1.3. Organization of the manuscript

This article is organized as following. The next section, section 2 is about our basic tools we have used, and explored. Section 3 is about our results and discussions, and how our work fit in on the scientific discussions of the openAI API new features, our challenges and what we did well. Finally, on section 4, we present our final words, we close the work by making a short description of what we have discussed and future works and how our work can be useful. We also provide a list of used references on the reference section.

Two Supplementary Materials are provided: one with the model training details, and the other with real conversation using the chatbot we designed.

2. Methods

We have used mainly two set of algorithms for creating SnakeCat: i) SnakeFace Pires and Dias Braga (2023); and ii) openAI APIs. SnakeFace is a transfer learning based set of models used to classify snakes focused on Brazil, it is a prototype, a concept testing. Whereas, openAI APIs is a set of APIs focused on large language models (LLMs), being chatGPT one of them.

On Section 2.1, it is explored the AI that has been used. We did our best effort to pass all the details necessary for reproducibility. In case of interest to reproduce, it would necessary learning more about the APIs before trying, we are assuming the reader knows the basic. A tutorial or just the official documentation must be used alongside.

On Section 2.2, parametrization of SnakeFace is discussed, and details considered important enough were reported.

On Section 2.3, it is presented further details, that could be essential to understand the paper.

Finally, on Section 2.4, links and external material are provided.

The reader is welcome to get in touch if any information is missing.

2.1. openAI APIs explored

See Fig. 4 for a schematic diagram on how those algorithms fit together.

2.1.1 *gpt-3-5-turbo-1106*

This is the latest version of the chatGPT API. It will be the default for chatGPT API in the future.

2.1.2 *gpt-4-106-preview*

This is their most advanced model when we wrote this paper. It is their API for GPT 4, released after chatGPT; with vision capability and more complex reasoning. Several researchers are finding that this version is indeed superior to chatGPT, e.g. Ghosh et al. (2023).

2.1.3 *gpt-4-vision-preview*

This model is the vision behind GPT 4. When GPT 4 was released, their vision capability was largely commented upon. It can interpret complex images and get meaning from them. As one example, we have also tested in collaboration with another author this algorithm applied to medical images (see section 3.4).

2.1.4 *gpt-4-vision-preview*

This model is the vision behind GPT 4. When GPT 4 was released, their vision capability was largely commented upon. It can interpret complex images and get meaning from them. As one example, we have also tested in collaboration with another author this algorithm applied to medical images (see section 3.4).

The image you've provided is a dental X-ray, also known as a radiograph. This type of image is used by dentists to evaluate the health of teeth and the surrounding structures that are not visible during a regular dental examination. In this X-ray, we can see several teeth with their roots. The image shows the internal structure of the teeth including the pulp chambers and canals. There's also visibility of the surrounding bone. Dentists examine these images for various purposes, such as checking for cavities between teeth, looking at tooth roots, assessing the health of the bony area around the tooth, and checking the status of developing teeth. However, without being a trained dental professional and without more context, it's not appropriate to comment on the specifics regarding the health of these teeth or diagnose any conditions.

gpt-4-vision-preview reading a medical X-ray.

Sadly, as we have discussed with a medical doctor also doing research on those tools, it cannot make a precise diagnosis. It is similar to snakes: it cannot pinpoint precisely the snake's species. What we have learnt is that coupling those models with a more precise model can create something quite advanced. They have released what is called GPTs. It can be trained to better read images; the authors discussed with another author training those GPTs, and the results are promising. Nonetheless, we are going to follow a different step: we are going to use an external library, namely, SnakeFace, which is based in TensorFlow.js.

Section	vision	model selector	final response
section 1	<i>gpt-4-vision-preview</i>	<i>gpt-3.5-turbo-1106</i>	<i>gpt-3.5-turbo-1106</i>
section 2	<i>gpt-4-vision-preview</i>	<i>gpt-4-1106-preview</i>	<i>gpt-3.5-turbo-1106</i>
section 3	<i>gpt-4-vision-preview</i>	<i>gpt-4-1106-preview</i>	<i>gpt-4-1106-preview</i>

Table 1. Different settings for SnakeFace using openAI APIs.

2.1.4. Table of results

We have tested combinations of the algorithms we just mentioned Table 2. The goal is presenting how the algorithm will possibly behavior with different configurations. This can be easily added as a configuration option from the SnakeChat, with their respective costs.

See Supplementary Material for more examples, using more configurations.

2.2. Parametrization of the algorithms

2.2.1. Parametrization of SnakeFace

We have used SnakeFace basically as reported on Pires and Dias Braga (2023); differences in behaviors may exist since we had to retrain the models when writing the paper about SnakeFace, but we have not retrained the models now, when developing SnakeChat; we have just uploaded locally the pretrained models. A pretrained model is a neural network model used after training; it is a model with knowledge. ChatGPT is a pretrained model, as it is MobileNet largely used for image identification.

The models used here are below. Those links are used to upload the pretrained models locally, this is SnakeFace when we mention it: a set of pretrained models trained to identify snakes. See Pires and Dias Braga (2023) for more details. We should stress that those models are prototypes. A production model should have much more snakes to be interesting enough.

1. Fake coral snakes (model 1): https://teachablemachine.withgoogle.com/models/W9_d1u4f/
2. Fake vs. true coral snakes model (model 2): <https://teachablemachine.withgoogle.com/models/9vw2MfLJw/>
3. Model with variety (model 3): <https://teachablemachine.withgoogle.com/models/Sc8mKQsO/>

Tip. if you click on the link, the model will open on the browser and you can input images.

We are going to use *gpt-3.5-turbo-1106* to automatically decide which models to call; later, we shall also try *outgpt-4-1106-preview*, the GPT 4 API (section 3.3). This is called function calling by the openAI team. This function is interesting since you give functions for the chat model to decide whether it should look for external information. For example, if you ask the temperature in a city, and the model has access to an API about temperature, it can call the API, get the temperature, and give back an educated response; instead of making it up as pure chatGPT famously do.

This approach of connecting chatGPT with powerful external tools has been largely explored. For example, the group behind Mathematica, a famous numerical software, used this approach as so chatGPT will give educated responses.

2.2.2. Parametrization of openAI APIs

The step by step of the algorithm is on Fig. 4. The parametrization on each step is as following.

1st step: getting a description from the image

Bellow is the section of the code that calls openAI for getting an image description. It is coded in Typescript and used in Angular. One point: the user can provide a human description as alternative. The AI generated description is triggered just when no description is provided by the user. It was also tested using MobileNet, which is a general-purpose computer vision model, for getting a general view of the image (see SM).

```
1 //It will return a promise, with a call to openAI
2 return await this.openai.chat.completions.create({
3   max_tokens:300,//size of reponse, in tokens
4   model: "gpt-4-vision-preview",//model used on the call
5   messages: [//Message sent
6     {
7       role: "user",
8       content: [
9         { type: "text",
10          text:
11            'There is a snake on this image.
12             Describe it.
13             Focus on the colors and patterns.
14             Ex. the snake is black, red, and white.
15             The colors are in strip patterns.
16             The black strips appear more easily than
17             the red ones.' },
18         {
19           type: "image_url",
20           image_url: {
21             "url": this.selectedImage,
22           },
23         },
24       ],
25     },
26   ],
27 }).then((response)=>{
28   //once the promised is finished
29   //getting just the response
30   return response.choices[0].message.content;
31 })
```

2nd step: getting which models from SnakeFace to call

The process of calling openAI API is the same as previously. The only change is the parametrization.

```
1   const options:any = {
2     model: "gpt-3.5-turbo-1106",
3     messages: messages,
4     tools: tools,
5     tool_choice: "auto",
6   };
```

gpt-3.5-turbo-1106 is their basic model, it has all the basic chatGPT capability, and with a large context (the size of possible inputs to the model, basic chatGPT is about 2 pages).

tools are the SnakeFace functions explained as so the model knows what each model does. Explanation given:

1. Fake coral snakes (model 1): Should be used only on snakes that look like coral snakes. This model was trained to identify fake coral snakes: snakes that look like coral snakes. The response is the snake with highest probability: name and the probability;
2. Fake vs. true coral snakes model (model 2): Should be used only on snakes that looks like coral snakes. This model was trained to identify both fake and true coral snakes: snakes that look like coral snakes. The response is the snake with highest probability: name and the probability;
3. Model with variety (model 3): This model was trained to identify several snakes. It is a generic model for classifying snakes from different species. The response is the snake with highest probability: name and the probability.

It is important to stress that those texts are called *prompt engineering*. Most of the success of using the openAI API is testing those texts. Changes

may happen. Generally, we try commands, then we change looking for desired behaviors. We are constantly changing those texts, looking for better behaviors from the API, it may change in the future.

The links to the models were given previously, and used to call them.

3rd step: gathering all the response, and getting a conclusion

This is the most tricky part. Several scenarios may happen. For instance, the SnakeFace prediction may not be in accordance with the openAI API suggestion. This is where those possible conflicting information may happen. We are constantly testing ways to handle those possible undesirable scenarios.

On this last part, we use again *gpt-35-turbo-1106* for creating a final response. On the final information sent to *gpt-35-turbo-1106*, we have a JSON response from SnakeFace.

```
1  this.openai.chat.completions.create({
2      model: "gpt-3.5-turbo-1106",
3      messages: messages,
4  }).then((secondResponse)=>{
5
6      this.message_chat.push(
7          'Em conclusão:
8          \n ${secondResponse.choices[0].message.content} ');
9      }); e
```

Please, find it here the complete method as Gist (a GitHub file of code).

2.3. Further details

2.3.1. Using Teachable Machine

SnakeFace was trained using Teachable Machine Pires and Dias Braga (2023).

We wanted to also experiment with models that have no overlapping: no overlapping between the species on the models. For the models originally published by Pires and Dias Braga (2023), there are overlapping between the models: model 2 contains model 1, and model 3 contains partially model 1 and model 2. On their scenario of research, it made sense such superposition.

For the case discussed herein, it becomes an issue. The same approach from Pires and Dias Braga (2023) was used for our models. Model 1 is the same, except we have retrained the model. We have discarded models 2 and 3 for having overlapping; but their results are still on the paper (Table 2). They were replaced with a model just for true coral snakes, and one model just for *bothros* genus. The species are the same, we have just reorganized them into separate sub-models, and those sub-models have no species in common. We wanted to investigate the trade-off between large models and specialized models on our chatbot.

The training graphs were added as Supplementary Material.

2.3.2. Model validation

The validation of this system is tricky. Validation is the process by which we make sure our system can achieve its goal, on a statistical meaningful level. One possible validation for our case is seeing how well it classifies snakes. The classification is done in three steps: i) get an image textual description (which is provided automatically by openAI API in case it is not provided by the user); ii) classify by image classification (SnakeFace); iii) use the information so far for making a final guess (openAI API).

The system is composed of two parts: SnakeFace and openAI API. Therefore, the first step is validating them. SnakeFace was validated on our other

paper Pires and Dias Braga (2023) and also on the SM, whereas the openAI API is validated on their respective official documentation, we provide links to the documentation when we use the models. Therefore, the validation of the parts are not our concern, it is out of scope. We assume they were properly done.

The second part that could be validation is how the system works as a whole, the sum of the parts: how the image classification provided by SnakeFace is supported by the openAI API, which is our goal. We believe it is meaningless this validation, we shall not do it. The reason is: a single change on either of the APIs we are using will make changes on the SnakeCat behavior: making the validation outdated. OpenAI APIs are constantly updated, and SnakeFace is being developed. Thus, the best we can do is trusting their respective validations.

Our validation is seeing how the system behavior, which we did at the discussion section. We have also provided a SM material for several real conversations with comments and comparing with other alternatives. That is the best validation we believe we can do, that makes sense for SnakeChat.

2.3.3. Inputting images

The current model is built to be used, by a non-expert user. The model is built on top of other models/APIs. Any image that could be uploaded to a browser will be just fine for classification; you can try to upload the image on our prototype, if the upload appears to you on the browser, it is fine. All the necessary image preprocessing is either done by us or by third-party APIs.



Figure 1. A *bothrops alternatus*, hard to classify. See that during the training, it was possible to classify this image, but as we added new images/species, we lost this good result. The goodness of the model is measured globally, not for a specific image or species. See Pires and Dias Braga (2023) for more details. Source: Pires and Dias Braga (2023).

The ideal image is a clear image. If you add a non-clear image (e.g., Fig. 1), the model may get confused, and misclassify; in some cases, the misclassification can be bad, like classifying *Bothrops alternatus* as *micurus corallinus*. See SM for another example of image that is hard to classify. Our hope is that GPT-4 with their description capability will support on this kind of mistakes. GPT-4 Vision can see and describe Fig. 1 properly. Interesting enough, SnakeFace, which uses transfer learning, cannot see very well colors. See SnakeFace for more details Pires and Dias Braga (2023), once this is the model that makes the image classification. Regarding the image description by GPT-4, see their documentation for more details.

2.3.4. Transfer learning

Our model is using models that are based on *transfer learning*. How those models are built is discussed on the Method section of Pires and Dias Braga (2023). On this aforementioned paper, it is possible to see the performance metrics for the models. The models that have been used herein were retrained, and most likely it will be retrained as we add new species, and run new experiments. Therefore, it is essential to focus on the qualitative behavior (e.g., whether the accuracy is higher enough and the loss function is low alongside the testing curves), not the quantitative (e.g., whether we get a predefined number on accuracy). Once the models are retrained, it is focused on the overall behavior. See Pires and Dias Braga (2023) for getting familiar with how those models are trained, and what to expect from those training procedures. On the case presented, it was used models already trained.

The chatbot that is discussed herein does not depend on the model used under the hood, neither it depends whether a specific technique is used (e.g., transfer learning). Therefore, it was made the decision not to focus on how the models under the hood were trained, it is discussed on Pires and Dias Braga (2023) and the SM. The reader is gently referred to this paper in case of interest on the models under the hood, which is out of focus of the current paper. Any model can be added under the hood, instead of the ones that were used. Any metrics presented herein would be most likely outdated soon since it is planned to add new species and test new approaches, which shall be discussed on future publications. Any new results will not affect the current paper, since the chatbot uses the models, but it does depend on those models for working.

One interesting fact about our approach: "the brain" (i.e., the models) is separated from the interface (i.e., the chatbot).

2.4. External resources: getting to know our prototype

In addition to the links already provided, you can also access a version of the model deployed here. You can upload an image and test it yourself.

Currently, we can classify these species, see SnakeFace for more details Pires and Dias Braga (2023). When we use the openAI API, we are shifting between those models. The task of the openAI API is creating a conversation, and call the proper model for answering the question; currently, you just upload the image, and the model will decide which model(s) to call, and add some possible context. If you upload an image outside of the list, the model will try to fit on this list: it will take the highest one.

Please, for more details, see Pires and Dias Braga (2023). Also, if you upload a no-snake image, it will still classify. We have tested on SM using MobileNet for avoiding these undesirable behavior. It can be solved adding a filter on the image workflow, which will avoid no-snake images entering the classification workflow; MobileNet can do it without any need of training. That something we intend to work on in the future, it is straightforward to implement (see SM for some initial testing).

2.4.1. model t: just fake coral snakes

Namely: 1) *Apostolepis assimilis*; 2) *Erythrolamprus aesculapi*; 3) *Oxyrhomblor*; 4) *Lamporellis triangulum triangulum*.

See SM, this was modified on the current format.



Figure 2. SnakeFace logo.

Lampropeltis triangulum triangulum was eliminated for not being concentrated on our geographical region of interest.

2.4.2. model 2: false coral snakes vs. true ones

On this model, we are using the following true coral snakes: 1) *Micurus corallinus*; 2) *Micurus frontalis*; 3) *Micurus alleni*; 4) *Micurus albicinctus*. For the false ones: 1) *Apostolepis assimilis*; 2) *Erythrolamprus aesculapii*; 3) *Oxyrhopus rhombifer*.

2.4.3 model 3: experimenting with diversity

For the fake coral snakes: 1) *Apostolepis assimilis*; 2) *Erythrolamprus aesculapii*; 3) *Oxyrhopus rhombifer*.

And for the diversity, we have used: 1) *Crotalus durissus*; 2) *Bothrops alternatus*; 3) *Bothrops neuwied*; 4) *Bothrops jararaca*. Since we already know the difficulties between false and true coral snakes, we left the true coral snakes out.

On the latest model, see SM, we have added also the true coral snake on this model.

3. Results and Discussion

At the end of our last paper Pires and Dias Braga (2023), we proposed an app that would connect SnakeFace with the textual capabilities of chatGPT (Fig.3).

SnakeFace is a transfer learning based tool created to classify snakes. Different from general-purpose computer vision models such as MobileNet and INaturalist Computer Vision, the model was designed to identify Brazilian species. General models, such as MobileNet, tend to overlook Brazilian species; they are good at saying whether we have a snake on the image, even adding a label, but bad at actually finding the species. As we are going to see, also the vision model of openAI also overlook Brazilian species. I have also compared with other models (e.g., BARD from Google, see Supplementary Material). Our goal is making SnakeFace a part of a conversational AI, and still getting it right for native species from Brazil. See that

BARD also can work as a general-purpose snake classification chatbot, see SM for a comparison.

If you see with attention, our first idealization of SnakeChat (Fig. 3) was by using chatGPT for guiding which algorithm should be used on a possible design with small models (models designed to identify snakes that have something in common, such as being false coral snake). We had in mind a manual function calling, now openAI made available an option to automatically do that. Furthermore, since we have already done a couple of tests, we knew chatGPT could get close to guessing snakes by textual descriptions. Then we thought: that could be useful. It is possible that on a real-accident with snakes, a textual description is the only information we may have for deciding on the serum.

The motivation on specialized models is that as we grow bigger, also, the accuracy per class will decline. When we have general-purpose computer vision models, such as MobileNet, it is impossible to make sure a species is well-classified. In the case of Naturalist, making a mistake will just imply that someone else will correct you. On a real application, such as finding snake serum, a mistake may be costly. Hence, the higher the accuracy per class, the best.

One interesting observation about snakes, compared to other animals. They are geographically-focused species. If you consider a healthcare center which may want to use the system, it could be severed with a local model. Trained and adapted for them. Thus, increasing the accuracy, and minimizing misclassifications. Even though some species of snake may have a broad geographical distribution, when we fix a geographical radius of interest, the possible species to be found on that region will drop drastically. Therefore, using general models will just increase the misclassifications, without a real trade-off at stake.

Small models are interesting due to two reasons mainly: i) when you retrain the model, we do not lose what we did good at the training stage, with a big model, when we retrain, we lose what we did well on the previous training; ii) small models may avoid mistakes between snakes that are very different.

The model does not see as we see, it seems to see globally, the surrounding can deceive the model to classify badly; one proof of that is that if you remove the background, surprisingly, it does not improve the prediction, it may worsen the prediction. Also, one of the snakes we tested (see Supplementary Material) was identified as a boat by MobileNet, and as a duck by INaturalist Vision Model. It makes sense since those two misclassifications are generally seen in water, and the snake was in water. Also, a very easy snake to classify was misclassified for being on a human hand (see SM).

One big issue we face with big models: as we add new species, we loose what we did right. For instance, during the first versions from SnakeFace, we could classify Fig. 1 and. Fig. 6. They are two problematic images: one is not clear, the other is a mimicking snake which does its work pretty well. The model during training is evaluated globally (a global error function is used to guide and finish the training process): it is not evaluated for a specific image. This is necessary since models are trained to operate on uncertainties, not for specific images

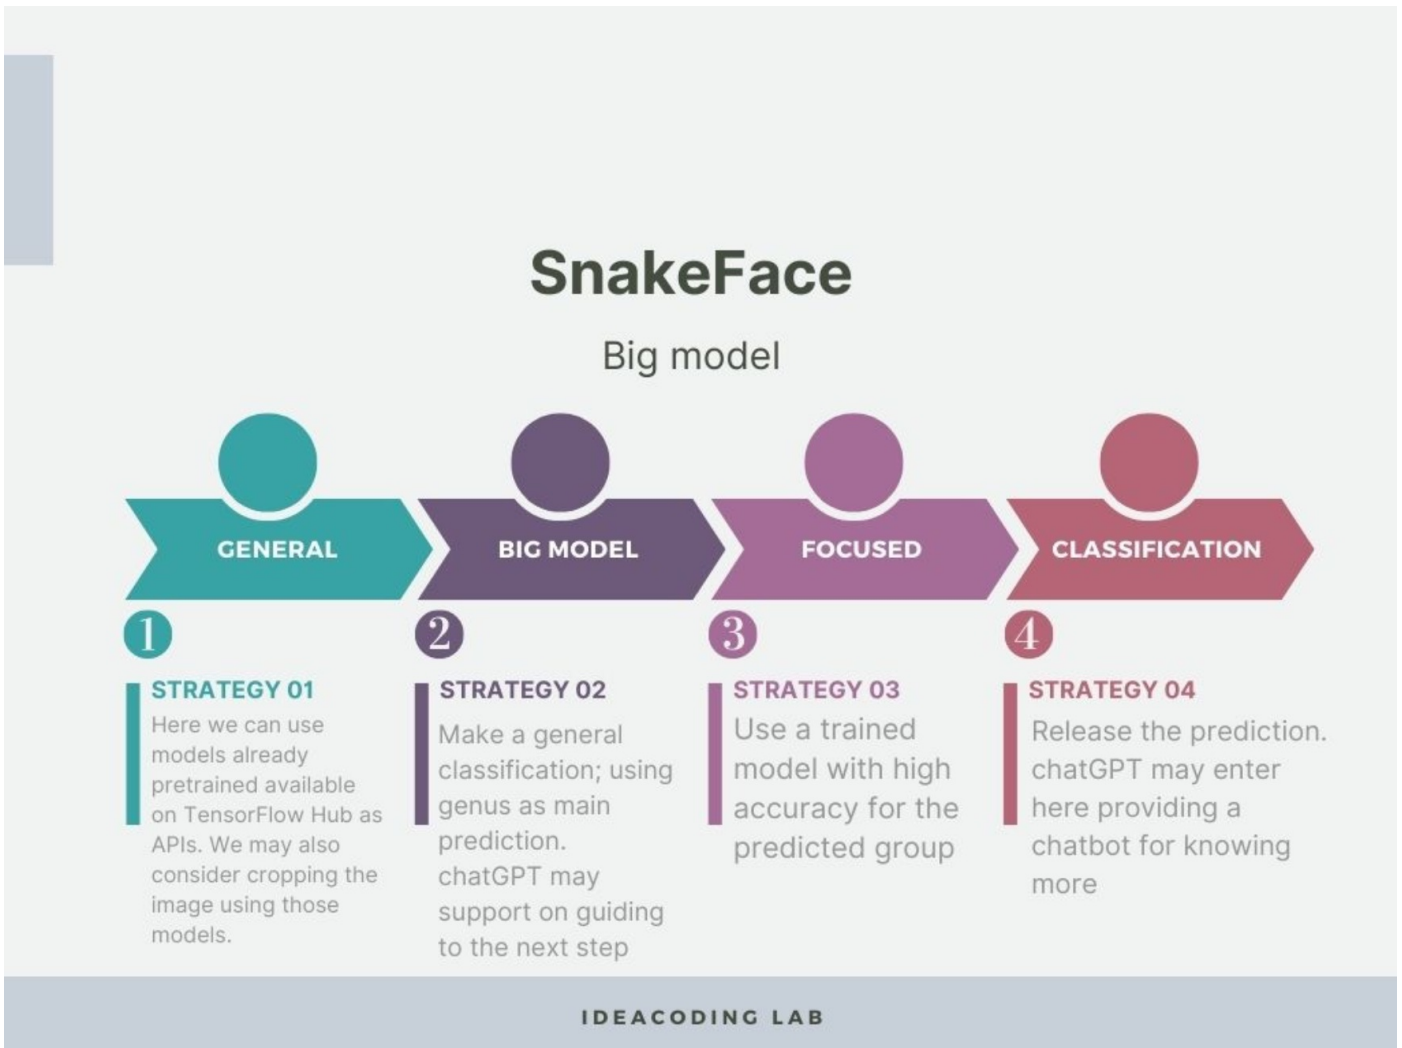


Figure 3. First idealization of snakeChat. Source: Pires and Dias Braga (2023).

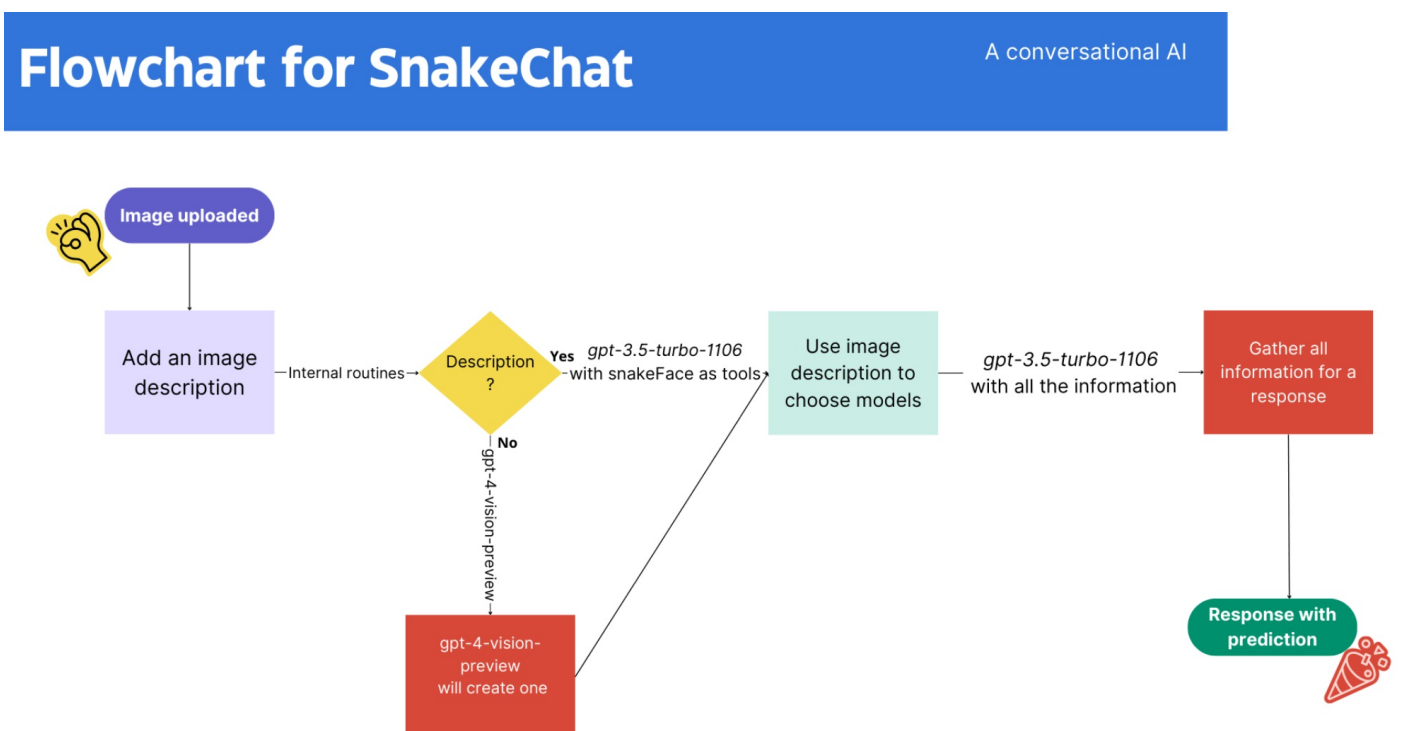


Figure 4. Executed version of SnakeChat based on SnakeFace vision. See that for the last test we present we have replaced *gpt-3.5-turbo-1106* by *gpt-4-1106-preview*. Legend: the

user uploads an image. If the user does not provide an image description, we use openAI API for creating an image textual description. Once we have the image description, we use openAI API to decide which model(s) from SnakeFace to use for making a guess using computer vision. Then, finally, we gather all the information (SnakeFace + openAI API) for making a final guess.

Recently, in a single announcement, openAI, the group behind chatGPT, released new tools for their APIs. Those new releases are what we need to make SnakeChat possible. Before, we were just thinking about chatGPT as just support, and informative for the user. Now, we are actually using their API for making our app smarter. Thus, it was integrated into SnakeFace.

On the following image (Fig. 4), we have how we executed fig. 3 using as pivot the latest APIs from openAI. Different from the initial plan to use just chatGPT API, we could go further thanks to the latest releases from openAI as APIs.

In the upcoming sections, we are going to discuss more closely three of the results presented on Table 2. The results on this table are for the first version of the chatbot, see SM for more variations of configurations.

See the supplementary material for more examples from the table.

It is important to highlight that this table does not present a statistical analysis. For doing so, we would need a considerable amount of work. openAI APIs are constantly updated, and SnakeFace is a prototype, new training will create new emergent behaviors on SnakeChat. SnakeChat is built on top of those models, thus, their behavior will change the behavior of SnakeChat. Thus, at the current moment, we consider useless any statistical analysis, which will entail also financial costs since the openAI APIs are not free of charge. Also, the latest configurations include more options. It is hard to predict how a LLMs will behave for each human-text input. We would need to test all the possibilities, variations and more. We also expect to add more features. The best approach is each user play with different settings, and see what works best for them.

3.1. Conversations

See SM for more conversations, using different configurations.

3.1.1. Oxorhous rhombifer

We are going to start with what the current app does quite well. Rarely, it will misidentify (Table 2).

First, we need to bring to attention, as so you can understand better what to expect and from where it may come some undesired behaviors (e.g., stochasticity).

- The algorithms from openAIs we are using are stochastic: they offer different responses upon the same input. This means that sometimes the response from openAI API may not be helpful as we hope. It seems to improve with their GPT 4 model, see our last conversation;
- SnakeFace is deterministic: the same input will always lead to the same output. This means that it may get it right, but be second-guessed by openAI API. This may occasionally lead to a confusing final response;
- We have not yet figured it out a good solution, but you can read the probabilities, and pick the highest one, but for some reasons, openAI API will pick not necessarily the highest probability as final response. It can be confusing. It seems to trust more the textual description than the numbers. It seems to improve with their GPT 4 model, see our last conversation (see lower sections from Table 2).

Species	SnakeFace	openAI API	Final result
<i>Oxyrhopus rhombifer</i>	got it right when using model for both true and false coral snakes	coral snake (wrong)	<i>Oxyrhopus rhombifer</i> (correct)
<i>Erythrolamprus aesculapii</i>	got it wrong since openAPI API called the wrong model	coral snake (wrong)	<i>Micrurus frontalis</i> (wrong)
<i>Oxyrhopus rhombifer</i>	got it right when using both models (just fake corals and both)	coral snake (wrong)	<i>Oxyrhopus rhombifer</i> (correct)
<i>Bothrops Jararaca</i>	got it right, openAPI API called the right model.	caninana (wrong)	<i>Bothrops Jararaca</i> (correct)
<i>Erythrolamprus aesculapii</i>	got it wrong since openAPI API called the wrong model	coral snake (wrong)	<i>Micrurus corallinus</i> (wrong)
<i>Oxyrhopus rhombifer</i>	got it right with two models, but misclassified on one model	coral snake (wrong)	<i>Micrurus corallinus</i> (wrong)
<i>Bothrops Jararaca</i>	got it right since openAPI API called the right model.	Jararaca (correct)	<i>Bothrops Jararaca</i> (correct)

Table 2. Summary of the results we are going to discuss.

Important. *this is not a statistical analysis, it is a general behavior demonstration.*

Legend. *The upper section is for gpt-3.5-turbo-1106 as final model (fourth column) and also as tool selector. Lower section is for gpt-4-1106-preview as final model, also as model selector. The second section (middle) is for gpt-4-1106-preview just on the guessing stage, for selecting which model to use from SnakeFace. See SM for examples of using gpt-4-1106-preview all the way. Note.* *The mistakes of SnakeFace (second column) is influenced when the wrong model is used, that is, it is influenced by openAI API (third column). The core idea of using openAI API was precisely to use the right model, and avoid the need of a big model. We believe that smaller models will give more control over the possible errors made by SnakeFace. openAI API can fix the error on a last call (third column), where it should gather all the information, and make a guess.*

It is important to bring to attention: variations on the responses are due to openAI APIs.

We have observed two common behaviors: i) when openAI API helps on the final response; ii) situations where it does not help, creating confusion.

One example is when SnakeFace gets it right, but openAI API guesses it wrong. Since we are logging all the responses, someone with the ability to handle those contractions may still gain from the app. Nonetheless, our hope would be an app that handles a single, simple, and correct response. It is not our current case. We saw apparent improvements when using their GPT 4 model: it does not call unnecessary models, therefore, less unnecessary numbers to confuse the final response. However, as you can see from Table 2, this behavior from GPT 4, API is not deterministic, it is not always the same. The lower sections all use GPT 4, API, as so we can have an idea on how it may be used to enhance the SnakeChat behavior.

Another example, see SM, is when a human-entered snake description was provided. openAI API called the wrong model. By calling the wrong model, SnakeFace provided a wrong prediction, which lead to a wrong final response. What is interesting: GPT 4, API created an entire argumentation to back up the wrong prediction, which was inconsistent with the snake description. It is a sort of "hallucination". I have entered the same conclusion on Bing Chat and Bard, they are both LLMs, asking for checking for inconsistency, they both failed to see that the snake description was inconsistent with the predicted snake. It is most likely that anyone minimally knowledgeable would see the inconstancy.

One nice feature from openAI API that may come handy in the future is that they allow to train their textual models, called fine-tuning; and then the model will be available to be called. It is possible in the future try to improve those textual guesses by training their model, and using the fine-tuned model instead of the generic one we are using. It is similar to what we did with SnakeFace: we fine-tuned a generic model called Mobilenet, that could not identify precisely Brazilian snakes. MobileNet was trained for generic objects. It is safe to assume that in some situations, a text-description of the snake is all someone can provide after an accident with the snake, which most likely will run away. Or a human-memory description.



Figure 5. *Oxyrhopus rhombifer*. *Oxyrhopus rhombifer*, the **Amazon false coral snake**, is a species of snake in the family Colubridae. The species is native to Brazil, Bolivia, Paraguay, Uruguay, and Argentina. Source: Wikipedia.

Our input image is below (Fig. 5): it is a clear image, no visual distractions, simple background. A good question is how the model handles bad images (e.g., Fig. 1): it is uncertain. It may see, or may not see. As SM, there is a case where it is hard even for humans to see the snake hidden on the vegetation, but the model saw, even though it did a misclassification with a similar snake. Other models, such as MobileNet, failed to see the snake at all. The ideal is always a clear image, with no obstructions.

Keep in mind: it is a matter of training, since the image mentioned was classified on previous training, but we lost it as we added new species. It is something we cannot control.

We are going to test first a human-generated image description. As SM, we have tested in PT since this is the final goal, and we want to make sure the chatbot can handle Portuguese properly.

The snake is black, red, and white. The colors are in strips. The strips black are bigger, whereas the red ones are smaller.

We asked *gpt-3.5-turbo-1106* guessing the snake from the description. It is a parameter for each function; at the current state, it is a dummy parameter, we are just using as chat logs. It guessed as *coral snake*, which is wrong. As said before, the behavior of openAI APIs is stochastic, in some tests, they actually guessed it right.

Response when calling model for both true and false coral snakes.

*Probabilidade de 82% de ser uma **Oxyrhopus rhombifer** usando SnakeFace, um sistema de identificação por imagens. [correct]*

Obs. the system was designed to respond in Portuguese, thus, you will see hybrid responses.

Second call, using the generic model:

*Probabilidade de 58% de ser uma **Micrurus corallinus** [this is a true coral snake] usando SnakeFace, um sistema de identificação por imagens.*

Remember that openAI API decides which model(s) to call: those calls were done automatically by openAI API. It has at its disposal three models, discussed in Pires and Dias Braga (2023); see the method section for minimal details therein. See SM where we have tested providing more models, and more configurations.

If we see logs, *gpt-3.5-turbo-1106* decided to call two models we made available from SnakeFace: *trueandfalsecorals* and *genericmodelsnake*; the former is used to classify false and true coral snakes, it was created to try to separate them from each other, see Pires and Dias Braga (2023); whereas the latter is for generic snakes, it was meant for a possible big model, generic. It did not call for the false coral snake model: which makes sense since its guess was wrong, it was coral snake. See that in other tests we have done, it actually guessed it was a false coral snake; remember that its behavior is stochastic.

Thus, the first call to SnakeFace rejected *gpt-3.5-turbo-1106* guess on the snake, whereas the second guess confirmed.

However, on this case it actually joined properly the contradicting information:

*The snake in the photo is most likely an **Oxyrhopus rhombifer**, also known as the Mehuaco or South American False Coral Snake. It is often confused with the **Micrurus corallinus**, or the Coral snake, which has similar coloration.*

This response was perfect: *gpt-3.5-turbo-1106* was able to actually handle the fact that the generic model and *gpt-3.5-turbo-1106* confused the snakes that are alike. The generic model of SnakeFace gave 15% probability for *Oxyrhopus rhombifer*. The false vs. true coral snake gave lower than 1% of being a *Micrurus corallinus*, which is good. It reinforces our call for specialized small models for SnakeFace instead of big models, as we are proposing herein.

If we make an analogy to human: *gpt-3.5-turbo-1106* on this case specifically did not fall into confirmation bias. It seems based on observations during the last

conversation using *gpt-4-1106-preview*: their GPT 4 model seems to be less likely to fall into confirmation bias, as we use it here (the model uses SnakeFace mistakes to confirm their initial guess).

One may ask ourselves how other researches are doing. We did a small literature survey on our other paper Pires and Dias Braga (2023), we are going to replicate it herein. In essence, there are several researchers also applying transfer learning to snake classification, it seems a topic that got researchers interested, and the results are in the same direction as we are going. Except, none of them is using chatbots, or even, JavaScript. Those two details may give us an advantage on building an app. Also, they seem to focus on big models. Big models can increase problems on accuracy per class (it may perform well on a set of classes, but badly on others). Low accuracy per class can become an issue when it is important to avoid misclassifications (e.g., finding a snake serum).

One issue we found is that those models are built using different models and they are in python, we are coding in JavaScript. None of those models made available, as far as we know, APIs available, that we could just plug in to our SnakeChat and test. We are opened for collaborations with other researchers to test their models inside SnakeChat. All the models used herein were used as legos: we can just remove one and put others for testing (see. Fig. 3). This approach makes it easier to compare models, as long as they are easily available for being used.

Now, we are going to test the system without human-image description: the AI (*gpt-4-vision-preview*) will make the description.

The snake in the image has a patterned coloration with three primary colors: black, red, and white. The black color is prominent and appears in thick bands or stripes that run the length of the snake's body. Between the black bands are patches of red color, which are themselves outlined with a thinner white or cream border that accentuates the red areas. This tricolor pattern is highly striking and gives the snake a distinct and bold appearance. The arrangement of colors suggests a repeating and systematic pattern, which is often characteristic of certain snake species, particularly those that use their distinct coloration as a warning or as camouflage.

The snake guess was still coral, even with its own image description. We could test fine-tuning their algorithm, something we may test in the future. For this approach, we need image description and their respective snakes. We are not even sure if that would be possible.

In that case, the final response was wrong, even though two of SnakeFace models predicted the correct snake (*Oxyrhopus rhombifer*) : 82% and 69% ; one of the models misclassified as *Micrurus corallinus* (58%)

But still, the final response was:

*The snake in the photo appears to be a coral snake, specifically a *Micrurus corallinus*. The tricolor pattern and the prominence of black, red, and white colors are characteristic of coral snakes.*

As we said before, someone reading all the logs may say: "it is two against one", and pick the correct guess. Making an analogy to human psychology: here their algorithm felt into confirmation bias.

We tried out their latest model, GPT 4 (*gpt-4-1106-preview*), and also got similar response:

Based on the description provided, the snake appears to resemble the coloration and pattern commonly seen in a Coral Snake or a species mimicking a Coral Snake, such as a Milk Snake or a King Snake. [that was a better guess!]

*Coral Snakes, which belong to the genus *Micrurus* or *Micruroides* in the New World, and the *Calliophis*, *Hemibungarus*, or *Sionmicrurus* genera in the Old World, are known for their distinctive red, black, and yellow/white banding. However, the phrase "red touch black, venom lack; red touch yellow, kill a fellow" is a mnemonic often used to distinguish the venomous Coral Snake from its non-venomous mimics in some regions (like the Milk Snake or the King Snake) where the order of the colors varies. Nonetheless, this rhyme is not reliable for identifying Coral Snakes globally, as color patterns can differ by region.*

Given that the description mentions red and white (or cream), and considering that some Coral Snake species may have a yellowish or whitish band instead of a clear yellow, the snake could potentially be a Coral Snake. On the other hand, since the description does not specify yellow but white or cream, it might be more likely the snake in question

What is interesting about this response, the algorithm gave the answer when it has written a common saying amongst people: *red touch black, venom lack; red touch yellow, kill a fellow*". There is no yellow on the snake coloration. It finished: *the description does not specify yellow but white or cream, it might be more likely the snake in question*. The conclusion was strange, it seemed it would actually conclude right.

But when we asked to make it short: Coral Snake. Same as the current version we are using. Thus, what remains is trying to fine tuning their model. We are going to leave that experiment for a possible future experiment.

3.2. *Erythrolamprus aesculapii*

As we brought to attention in our previous work Pires and Dias Braga (2023): this snake is tough to classify since it was able to fool our models most of the time, mimicking a coral snake. One thing we said in our previous paper: our goal should be avoiding identifying venomous snake as non-venomous. We called it error type II. In fact, it was the main focus of another research Felipe Guimaraes Dos Santos (2021). Since we are using SnakeFace, we did not expect different! See that we were able during the training to classify it properly, this image. The problem is that the model is

evaluated as good and bad globally, not for a specific image: when you see a good result from the image, it is an overall model behaviour. The version we have used now was retrained after the paper, and it fooled the current model.

We added the description manually for this case:

The snake is red, white and black. The colors are in strips. The red strip is the biggest ones, whereas the black and white ones seems at the same size.

We asked also their latest model to guess, it guess equally wrong compared to the current version we are using from their API. Guessed coral snake (see that this behavior is persistent for coral snakes, see Table 2).

As a result, it decided to use just one model. Guessed as *Micrurus frontalis* (100%); indeed, compared. Fig. 6 and Fig. 7, it is hard to make the difference just using images. We believe that if the user provides details that are not in the image, we believe the whole system (i.e., SnakeChat) can get it right. The issue with that is that just an expert would think to provide information about the head, tail and more. We have in mind a system for non-experts. We have experimented with asking GPT 4 Vision to focus on this details when generating an image description, and the result is interesting (see SM). When those details are visible in the image, it can describe those parts of the snake, even for images as bad as Fig. 1.

3.3. Bothrops Jararaca

We are going to try something different; we tested as we are, and the guesses were bad, calling *Bothrops jararaca* a coral snake. It may lead to right answer corrected by SnakeFace and accepted, but also to what we have already seen: a sort of confirmation bias. See lower sections from Table 2.

So far, we used their cheapest model, let's try out GPT 4 (*gpt-4-1106-preview*). This is their latest model, and widely commented upon, with intriguing results like reading images and graphs with high level of intelligence.



Figure 6. *Erythrolamprus aesculapii*. Source: Pires and Dias Braga (2023).



Figure 7. *Micrurus frontalis*; compare with fig. 6 , they are essentially the same visually. Source: Wikipedia.



Figure 8. *Bothrops jararaca*. Source: Wikipedia.

Description using the same algorithm we have already used so far:

The snake in the image has a patterned appearance with a mix of neutral and earthy tones. Its scales seem to have a sequence of dark brown or black, beige, and lighter brown hues. The patterns are irregular, with some semblance of bands or blotches that alternate along the length of its body. The darker bands are bolder and more pronounced, providing a type of camouflage among the twigs and leaves on the ground. There do not appear to be any red or white stripes as part of this snake's coloration.

Surprise as it may seem, this description left to guessing as coral snake with the model used so far. We are changing both the function calling algorithm and the final response algorithm (see 3; see also Table 2, lower section). Thus, just the image description algorithm was kept; at the moment, it is the only option they provide to describe images.

Their guess was jararaca. Much better. We have done the upload three times, with the same result. We found that the *gpt-4-1106-preview* is more predictable than *gpt-3.5-turbo-1106*. In their website, they suggest to fine-tune *gpt-3.5-turbo-1106* for getting this better result, just using the outputs of *gpt-4-1106-preview*. *gpt-4-1106-preview* is more expensive to use. Regarding the final cost of the fine-tuned model, it does not seem a good deal once the fine-tuned model will be more expensive than using their GPT 4 already trained. It is something we must consider if needed.

Our model confirmed the textual guess *Bothrops jararaca* (97%).

What is interesting, it also avoided calling unnecessary functions. When guessing as coral snake, as did the previous setting-up, it will call unnecessary functions, creating noise in the final response: creating space for wrong guessing the snake. On the calls we did with this image, they always called just one model: the general model; different from the previous set up, that seems to be calling unnecessary functions, and creating unnecessary results to interpreter.

Final response:

*Based on the description provided, the snake in the photo is likely to be a *Bothrops jararaca*, commonly known as the jararaca. It is a species of pit viper found in South America, and its appearance, characterized by a patterned skin with a sequence of dark brown or black, beige, and lighter brown hues, matches the description given. This type of coloration allows the snake to blend into the forest floor, making it well-camouflaged among leaves and twigs. Please note that without the actual photo, this identification should be considered a hypothesis and should be verified with a visual confirmation or by a herpetologist.*

Best of all, this result tends to repeat itself if you upload the same image. It is good! See Table 2 for a case where it guessed wrong, but still we got it right.

3.4. How SnakeChat could be adapted to similar applications

Even though we have focused on snakes, this approach can be applied to any case using images. The final result will be the same: a chatbot with state-of-the-art capabilities using images, a *conversational AI*. This includes, as we stressed on our other paper Pires and Dias Braga (2023), other species beyond snakes. As long as you can separate the species from the images, you can build a similar system for your desired application.

For instance, we have done some tests and discussions with a medical doctor researcher working on those applications as well. You can either train what they call GPTs (which is their way to make their models even more accessible to no programmers), or use the same path we used here:

Teachable Machine, and use the models as we did, as external APIs. We have tested Teachable Machine applied to medical images, and the results are promising. In fact, our next paper may be on a system we are working on, but for medical images instead of Snakes. Thanks to how the openAI APIs are built, the transition between applications is straightforward. We have done a first prototype here Pires (2023c).

Our path has the advantage that Teachable Machine is free of charge to use, all the way, including hosting the model on Google. This means you do not have to pay for the vision part. Also, the GPTs are not available as APIs, when we wrote this paper. The computer vision model we are using (namely, *gpt-4-vision-preview*) does not allow fine-tuning, as they pinpoint on their own documentation. We have also here a comment on their forum: it may be released in the upcoming feature as an API.

Another option is that you have a model from your group, which was well-trained for a specific task. Even though the models from openAI are advanced, they seem to perform poorly on specialized problems. In a couple of tests done by another researcher, and confirmed by us in our own scenarios, if you ask their computer vision models to make diagnoses from images, it will either make a wrong diagnosis, or refuse to do so Pires (2023c). But they are very good at describing images, as we did. This description can be used to choose more specialized models, as we did.

There are plenty of literature studying the openAI models applied to diagnosis in medicine, their limitations and strengths, as example, but not exhaustive. Ghosn et al. (2023) compared the GPT 3.5 and GPT 4.0. As we expected, GPT 4 was better at answering questions from an standard examination in radiology, and also explaining the answers.

3.5. Challenges

This section was created based on a public review we received on Qeios.com, on the first draft of this paper. We have found very insightful the challenges the reviewer raised, and found it interesting to add those challenges and their respective discussions on this final manuscript. Hopefully, it will support others on understanding our challenges.

3.5.1. Potential for disinformation

There is a risk for disinformation on those LLMs. One example is on the SM, where the LLM created an entire argumentation around a wrong prediction. That is a topic that we may need to handle on the future. Several solutions may be tested. As example, it is possible to provide texts for the models, and ask the model to base their answer on this material. That would require a dataset of texts written by experts, which may be a long-term task. Other option is experimenting with fine-tuning their models. Those are all possible ways to handle disinformation on those chatbots.

3.5.2. Hallucination

Hallucination is largely associated to LLMs, it can be also seen on computer vision models. For instance, MobileNet saw a boat where we a snake, or even a crocodile. See SM. This is something we must consider solutions to mitigate. One way is to avoid using the wrong model, even though we have tested some possibilities, such as using MobileNet for having a general classifications, I am planning to also experiment with other alternatives (e.g., Naturalist).

3.5.3. Integration complexity/dependency on External APIs

Integrating external data sources (APIs), like Naturalist and openAI APIs, as we did and are planning to do, adds complexity to the system. Relying on external APIs, such as OpenAI and INaturalist or any other we may add in the future, introduces a dependency on the availability and stability of these services. Any disruptions or changes in these APIs could impact the functionality of the system. That is indeed a possibility when our system is built using external parts, third-party software/APIs/datasets. It is a trade-off: speed at building comes at a cost, lack of control of the system behavior. Moreover, it is hard to imagine we could build a dataset so rich as Naturalist, where people are posting their observations constantly, they even have a pretrained model on TensorFlow Hub Pires and Dias Braga (2023). In our view, this sharing-culture we built in the last decades is the best thing that ever happened to applied sciences, to collaborative endeavours as ours.

In fact, that was the detail that made it senseless a validation, as we did with SnakeFace Pires and Dias Braga (2023): we plotted errors graphs, calculated confusion matrices and more; we did a standard validation of neural network models, following well-accepted metrics and procedures. A single change on any of those external sources (i.e., APIs) will change the validity of the validation.

Systems like ours nowadays fall into the complex system category, a theory largely used on studying biological systems. Complex is not about size, when seen from the complex system theory, it is about how hard it is to predict behaviors, as we like to say "the whole is greater than the sum of its parts", first coined by the philosopher Aristotle. In fact, we have used a term from complex systems theory on this paper, we have used the term 'emergent property' for designating SnakeChat behavior, it is the summation of SnakeFace and OpenAI APIs working synergically, which by themselves depends on other parts that we cannot see and predict; SnakeFace is more predictable than openAI APIs for being a deterministic model.

One big criticism on openAI is the fact that their algorithms are closed, no one really knows what is inside the magic hat. They keep a page, and an e-mail list, where they constantly post instabilities, and what they are doing to handle the issues. It is constantly active: it seems they are always making changes in their models, which may create undesirable behaviors. For instance, when we simulated the results on Table 2, we had the impression that GPT 4, API was more predictable, but now it seems it is behaving as chatGPT API, just an impression.

A term that is gaining attention is *software entropy*. This term is used on software testing Khorikov (2020). We are using Angular as framework for SnakeChat, Angular was created to be tested Jesse Palmer (2018). Software testing is done precisely to handle, or at least, mitigate software undesirable behaviors, making them more predictable. External APIs and datasets is a big challenge. We do what is called mocks: we just ignore them. We create a local version of APIs and datasets, which behaviors in a predictable way, which is unrealistic. The idea is that each groups should be responsible for their quality control. Therefore, generally, we do not try to control external APIs/datasets. The best we can do is choosing reliable sources. At the current stage, we are not concerned about software testing, we are prototyping. There are whole organizations just focused on software testings.

As we can see from Table 2, the behavior of openAI APIs is indeed stochastic, and it seems they are working to increase predictability of their APIs, which for sure will affect ours.

When we do a scientific search for terms like "chatbots in biology", most of the results, when we find any, are medicine driven, and they are using precisely openAI solutions. It is impossible to imagine someone building a whole large language model, just to classify snakes. The technology behind openAI was not created by them, it was around and it was created by scholars and Google researchers. Therefore, the fact that no scientific group created such a LLM is a sign that it is unfeasible considering creating our own LLM. The costs of openAI is very low, based on what it can do. Thus, it is hard to imagine other groups would build such a system with such a low cost. Of course, we do have open sources groups working on that, and their results are interesting. Hopefully, they will also create LLMs, like openAI APIs: easy to use, with almost zero setting ups.

3.5.4. User Understanding and Expectations

This is an interesting point, and we have discussed it previously in Pires (2022). We believe, similar to the mindset from openAI: we learn by interacting with the user. That is why we have launched all those prototypes. It is hard to handle user's expectation when you are using anything related to chatGPT, and we have seen that first-hand: all the attention around chatGPT can backfire, people seem to expect much more than chatGPT-related systems can delivery. Understanding the user may take time, but we see a general user, without expertise in biology. We also see

healthcare centers, and teaching institutions as final users.

3.5.5. Scalability Challenges

We are not at this stage, but certainly scalability is something to consider, it is largely studied on start-ups contexts Blank and Dorf (2012). One challenge of scalability is adding new species, there are more than 2.000 worldwide, and the best model we know was able to classify about 700 species; we discuss that on our other paper Pires and Dias Braga (2023). One solution is on this paper: work with small models, and use openAI API to handle them intelligently. When we add new species, we cannot know what will happen with the current results. For instance, on early versions of the general model from SnakeFace, we could actually classify both Fig. 1 and Fig. 6, two hard images to classify. As we added new species, we lost this impressive result.

Another challenge, which is discussed in Pires and Dias Braga (2023), is annotation, a good dataset. For training, we need images and their respective tags. Which have partially been solved using Google Image. Nonetheless, that is something to pay attention to. Our latest models were built using INaturalist, which is much better. They have a rich community, and annotations are under consensus. Also, their images are real, they are taken by users on real-scenarios. This is interesting for training the model.

One details in our favor: big models tends to lose accuracy per class. We can either build chatbots for regions (e.g., healthcare centers) or change the models according to the user geographical position. Also, we are considering a model that could better be used to chose small models. Small models could make the scalability easier.

3.5.6. Pricing

It is always nice to build such a app for free, or at reduced cost. Nonetheless, there are costs involved that will have to be taken into account on a production model. OpenAI API has a pricing that changes according to the model used. Their pricing is "pay-as-you-go": they charge as you use their models, not a monthly/annual billing. It is hard to make an estimation now since the pricing will change on which model is used. For instance, the GPT 4 Vision charges high resolution image higher. Since we are using more than one model at once, depending on the configuration (see SM), the pricing will vary according to the configuration chosen by the user.

One solution for creating a free version is replacing the function calling by MobileNet, which I have shown not to do a very good job with Brazilian snakes. It will work, but with less accuracy. I have also showed that it can be used instead of the image description model from openAI. Again, the result may not be good as it is now, see SM.

There are several open source LLMs gaining momentum as the LLMs gain attention for their advanced capabilities, and they may become in the future an option. I need to get more familiar with them, and see if it is possible to replace the openAI LLMs with them, whenever possible for reducing costs with LLMs.

There is also the server cost, I am using Heroku. Their pricing is for the entire server, it is possible to deploy several apps on the same plan. For this, options are Wix Site, which allows to code in JavaScript, and charge for premium options only. Another option is GitHub Pages. Of course, all those options come with limitations. More options may certainly appear as we go.

3.5.7. Limited Control over Training Data

It is true that we have limited control over what openAI API says. One solution is using a text-dataset on snakes we can classify, created by experts. It can be just a blog, since we can access blog posts and send the content to openAI API to use as reference to their responses. The issue with this strategy is precisely having this expert-curated text dataset. Then, instead of using the knowledge from openAI API models, we just use it as a knowledge-mining tool. The latest APIs from openAI now can read almost 300 pages at once, they increased considerably their "attention window", a point of large criticism, and also launched Assistants, this is an way to allow them to actually manage attention-windows automatically. This is something we have been exploring largely on other applications where we use chatbots. See a discussion we did on data science Pires (2023a).

3.5.8. Adding more species

According to *Encyclopaedia Britannica*, there are nearly 3,000 species of snakes distributed nearly worldwide. Brazil is home to a diverse range of snake species, with nearly 400 species found in the country according to Atlas of Brazilian Snakes. The good question is whether we could scale up. The scaling up is not a issue of SnakeChat, which is a chatbot, it is an issue of SnakeFace Pires and Dias Braga (2023), the snake classifier under the hood when the chatbot tries to guess a snake.

Machine intelligence is different from humans': it does not change very much as you add more classes. Humans have a "normal intelligence" (a reference to the normal distribution), it is stronger around expertise, where machines have a "flatter intelligence", they tend to be good at a big number of classes at once. It is true that human experts can be better than machine learning in some scenarios, nonetheless, machine learning models once trained properly can be good at several cases at once. For classifying several species, one would need several experts. The expert will be good at their expertise, but average on the remaining. Machine learning once trained properly can be similar to several experts at once.

Even though Pires and Dias Braga (2023) explored a small number of species, it was explored strategically: similar snakes and diverse snakes. Those are like extremes. The current paper actually says that we do not need a huge model, but a number of small specialized models, and allow the chatbot to pick the proper model.

Thus, adding more species is more an issue of creating more models. Of course, this is time and resource -consuming, and challenges may appear on the way. See Pires and Dias Braga (2023) for more discussions. Scaling this prototype to cover a broader range of snake species will be challenging and will require significantly more data and resources. Nonetheless, this is not an impossible task, we have shown it is possible.

3.5.9. Gpt-4 API outperformed chatGPT API

Recent literature has highlighted the apparent superiority of GPT-4 over its predecessor GPT 3-5, particularly in the context of chat-based language models. Lubiana et al. (2023) provides tips for optimizing workflows using ChatGPT, which is powered by GPT-3-5 and GPT-4. Rahaman et al. (2023) compares the performance of GPT-4 with previous versions, showing improvements in training data, computation speed, answer accuracy, and overall performance. Rosol et al. (2023) reveals that GPT-4 outperforms GPT-3-5 in a medical examination, with higher accuracy. These findings highlight the improved performance and capabilities of GPT-4 compared to GPT-3-5, supporting the claim of its superiority in the literature Lubiana et al. (2023); Rahaman et al. (2023); Rosol et al. (2023).

We have mentioned regarding an apparent superiority of GPT-4 over chatGPT (i.e., GPT 3-5) for our case. It has not been done a careful statistical analysis, therefore, it is not safe for us making any hold claim on that direction. What we can say: the first simulations showed that GPT 4 tended to call just the needed model, which is a good behavior; but on the final simulations, which was done days later, they seemed to be equally in term of function calling. Calling wrong models as we have seen creates noise, unnecessary information, which can induces the final stage of the chatbot to arrive to wrong conclusion about the snake species. We also cannot say for sure why, nonetheless, we have noticed on another study we have done, which may be published regarding Section 3.4, is that the GPT 4 API is more "obedient", its higher cognitive ability makes it better at deciding which function to call. Thus, the function calling capability may work better with GPT 4 API. In fact, it was found that GPT 4 has a higher cognitive capability compared to GPT 3-5 West (2023).

3.5.10. Application on healthcare centers

One application I am considering is for the local healthcare center from Antonio Pereira (Ouro Preto). Snakes are local animals. There are groups of studies on the local university (Federal University of Ouro Preto) to register local snakes; SnakeFace Pires and Dias Braga (2023) was created in a collaboration with a student from this group. Since accident with snakes are real, the high-precision for a small number of snakes is good. Unless someone is by purpose entering a foreign snake, the number of snakes on a geographical location is small. SnakeFace/SnakeChat is good at a small number of snakes, different from general models which will misclassify those snakes (see TM). The type of antivenon serum required for treatment depends on the species of snake that caused the bite, and timing may be crucial.

3.5.11. Limitations and future works

In addition to all the testing I have done as part of the development process, I have also provided a set of conversation samples as Supplementary

Material. SnakeChat/SnakeFace is not perfect, and it has several flaws that we hope to improve with time.

For instance, as I was finishing the writing of this manuscript version, I have found a snake on the same geographical region the species I have added was, during my walks. I knew the species was not on the dataset. After entering the image, the system have chosen a snake that had similar color, but a very different species. This happens because by prompt engineering on the openAI API, I have directly instructed to use the models even when the snake identified by MobieNet/GPT 4 Vision was different. This was done since initially the function calling would not call when the snake on the image was different from the ones on the models. As an example, MobileNet will identify *bothrops alternatus* as a rattlesnake, which are very different snakes. The function calling API would refuse to use any model since they are very different snakes.

Thus, we have a tradeoff, which seems without solution so far: a model was created since the general models we found currently online cannot see the Brazilian snakes well, especially, from the geographical region I have concentrated. However, when a wrong snake is entered, the model would still classify running the risk of making a totally wrong classification. I have actually tried BARD for the snake mentioned, and it did an even worse prediction. The solution was a manual combination of INaturalist with Google Image, which I have done manually, and took me minutes to merge the pieces of information. Therefore, our system can see what it was trained to see on where it is, but it can see what it was trained to see what it does not exist (i.e., a snake on wrong models).

This is a general issue on computer vision models, not specific for our case. If you enter an image with just grass and no snake, I have done that as a placebo, see SM, MobileNet saw several animals, including a rabbit. For a bad image, see SM, it saw even a crocodile. In our case, since the system is specialized, it can become easier to see those "hallucinations". I am considering solutions, such as using other general models rather than MobileNet, or GPT Vision. Nonetheless, it is important to keep in mind that some of the issues are bigger than our system, it is a computer vision issue, a general issue. There is no artificial general intelligence currently, all AI nowadays are designed to solve specific problems, as I have done. Some are more specialized (e.g., SnakeFace), other are more generic (e.g., BARD).

4. Conclusion

On this article, we have discussed a prototype for a conversational artificial intelligence (CAI) focused on snakes classifications. We have used the APIs from openAI as leverage, as so we did not have to build a CAI, just use it as leverage. Our model is an extension of SnakeFace, a transfer learning based model built to identify snakes. SnakeChat is focused on proving a more human-like prediction: instead of just providing a species and probability, we provide as a textual response. It is a task-oriented CAI, but thanks to how openAI APIs are designed, shifting to a hybrid CAI is not that hard, also becoming a chat-oriented CAI.

In addition, we have used textual descriptions, that can be entered by the user or is generated by an artificial intelligence. Our simulations show a good integration between SnakeFace and OpenAI APIs, what we have named SnakeChat. Our next step would be: i) adding more species of snakes; ii) adding more information as so to enhance the predictions; iii) keep testing our model and keeping finding new algorithms that can make our system unique.

One way to add new information to the prediction is using INaturalist: they provide an API that can give back information about snakes, such as if they were ever seen on a given location. Several pieces of information, such as textual descriptions, probabilities, and accuracy on a geographical location, can enhance predictions. We believe that GPT 4 can integrate those information into a final response. One path is providing also INaturalist as possible functions to call, something we may test in the future.

We see a platform full of potential, especially, after all the advances in artificial intelligence in the last years, both in computer vision and conversational artificial intelligence. The best advances in AI was becoming public APIs, as so they became bricks for programmers, instead of final tools delivered already designed.

Supplementary Material

This material is available from the Supplementary data section and can be downloaded from the following links:

- snakechat-sm-2.pdf ([here](#))
- snakechat-v4-sm.pdf ([here](#))

Acknowledgments

This paper was reviewed on Qeios.com, an open-review process. It was more than twenty reviewers. Those reviews were very important for the new version of the paper, they gave me insights on how to improve. Let me thank some of the reviewers publicly. Namely: Atheer Alrammahi (University of Al-Qadisiyah); Andi Adriansyah (Universitas Mercu Buana); Muhammad Amir Malik (International Islamic University, Islamabad); Shilpa Nahajan (Dr. B. R. Ambedkar National Institute of Technology Jalandhar); Chirag Joshi; Elakkiya R (Birla Institute of Technology and Science Pilani); Degan Zhang; Venkatesan Venkatesan (Bannari Amman Institute of Technology); Rodrigo Marques de Figueiredo (Universidade Vale do Rio dos Sinos); Muna Hadi Saleh (University of Baghdad); Mei Hua Hsu (Chang Gung University of Science and Technology); Newlin Shebiah Russel *et al.*

References

- Blank, S. and Dorf, B. (2012). *The Startup Owner's Manual: The Step-By-Step Guide for Building a Great Company*, K & S Ranch.
- Desingu, K., Mirunalini, P. and Kumar, J. (2021). Snake species classification using transfer learning technique, *Conference and Labs of the Evaluation Forum*. **URL:** <https://api.semanticscholar.org/CorpusID:237298537>
- Durso AM, Moorthy GK, M. S. B. I. S. M. and de Castañeda R, R. (2021). Supervised learning computer vision benchmark for snake species identification from photographs: Implications for herpetology and global health, *Frontiers in Artificial Intelligence* **4**. <https://doi.org/10.3389/frai.2021.582110>
- Emsley, R. (2023). Chatgpt: these are not hallucinations – they're fabrications and falsifications, *Schizophrenia* **9**(52): 1–2.
- Felipe Guimarães Dos Santos, Emanuel A. Machiaveli, J. M. V. M. (2021). Snake classifier: Aplicativo mobile para classificação de serpentes peçonhentas. https://www.academia.edu/95702218/Snake_Classifier_Aplicativo_mobile_para_classifica%C3%A7%C3%A3o_de_serpentes_pe%C3%A7onhentas
- Freed, A. (2021). *Conversational AI: Chatbots that work*, Manning. **URL:** <https://www.amazon.com/Conversational-AI-Chatbots-that-work/dp/1617298832>
- Ghosn, Y., Sardouk, O. E., Jabbour, Y., Jrad, M., Kamareddine, M. H., Abbas, N., Saade, C. and Ghanem, A. A. (2023). Chatgpt 4 versus chatgpt 3.5 on the final frcr part a sample questions. assessing performance and accuracy of explanations, *medRxiv*. **URL:** <https://www.medrxiv.org/content/early/2023/09/08/2023.09.06.23295144>
- Jesse Palmer, Corinna Cohn, M. G. C. N. (2018). *Testing Angular Applications*, Manning Publications. **URL:** <https://www.manning.com/books/testing-angular-applications>
- Kalinathan, L., Balasundaram1, P., Ganesh, P., Bathala, S. S. and Mukesh, R. K. (2021). Automatic snake classification using deep learning algorithm, *Conference and Labs of the Evaluation Forum*, . <https://ceur-ws.org/Vol-2936/paper-135.pdf>
- Khorikov, V. (2020). *Unit Testing Principles, Practices, and Patterns*, Manning Publications.
- Krishnan, M. G. (2020). Impact of pretrained networks for snake species classification, *Conference and Labs of the Evaluation Forum*. **URL:** <https://api.semanticscholar.org/CorpusID:225073809>
- Kublik, S. and Saboo, S. (2022). *GPT-3: Building Innovative NLP Products Using Large Language Models* illustrated edn, O'Reilly Media, Incorporated.
- Lakshmi, D., chandra Panda, R., Amrita and Prakash, A. (2021). Life-saving app: Snake classification 'venomous and non-venomous' using fast.ai based on indian species, *Artificial Intelligence Systems and the Internet of Things in the Digital Era* **URL:** <https://api.semanticscholar.org/CorpusID:236671950>
- Lubiana, T., Lopes, R., Medeiros, P., Silva, J. C., Gonçalves, A. N. A., Maracaja-Coutinho, V. and Nakaya, H. T. I. (2023). Ten quick tips for harnessing the power of chatgpt/gpt-4 in computational biology. **URL:** <https://api.semanticscholar.org/CorpusID:257805086>
- Mirunalini, P., Desingu, K., Bharathi, H., Chodisetty, E. A. and Bhaskar, A. (2022). Deep learning and gradient boosting ensembles for

- classification of snake species, *Conference and Labs of the Evaluation Forum*. URL: <https://api.semanticscholar.org/CorpusID:251470924>
- Pires, J. G. (2021). An informal survey presents the gap between computer and medical doctors and biologists, Medium. Accessed on 6 December 2023. URL: <https://medium.com/theoretical-and-mathematical-biology/an-informal-survey-presents-the-gap-between-computer-and-medical-doctors-and-biologists-ca8816051739>
 - Pires, J. G. (2022). Innovating with biomathematics: the challenge of building user-friendly interfaces for computational biology *Academia Letters* p. Article 5792. URL: <https://doi.org/10.20935/AL5792>
 - Pires, J. G. (2023a). Data science using openai: testing their new capabilities focused on data science, *Qeios*. URL: <https://www.qeios.com/read/76QMHB>
 - Pires, J. G. (2023b). Do jumento ao einstein chatty: usando chatgpt para criar um assistente virtual inteligente usando a plataforma wix, <https://www.amazon.com.br/Jumento-Einstein-Chatty-assistente-inteligente-ebook/dp/B0C1P811QE>
 - Pires, J. G. (2023c). Robodoc: a conversational-ai based app for medical conversations, *medRxiv* **2023**: 2023.12– 31.23300681. URL: <https://www.medrxiv.org/content/10.1101/2023.12.31.23300681v1>
 - Pires, J. G. and Dias Braga, L. H. (2023). Snakeface: a transfer learning based app for snake classification, *Revista Brasileira de Computação Aplicada* **15**(3): 80–95. URL: <https://seer.upf.br/index.php/rbca/article/view/15028>
 - Progga, N. I., Rezoana, N., Hossain, M. S., ul Islam, R. and Andersson, K. (2021). A cnn based model for venomous and non-venomous snake classification, *Analogical and Inductive Inference*. URL: <https://api.semanticscholar.org/CorpusID:236963195>
 - Rahaman, M. S., Ahsan, M. M. T., Anjum, N., Terano, H. J. R. and Rahman, M. M. (2023). From chatgpt-3 to gpt-4: A significant advancement in ai-driven nlp tools, *Journal of Engineering and Emerging Technologies* URL: <https://api.semanticscholar.org/CorpusID:258650483>
 - Rajabizadeh, M. and Rezghi, M. (2021). A comparative study on image-based snake identification using machine learning, *Scientific Reports* **11**(1). <https://doi.org/10.1038/s41598-021-96031-1>.
 - Rosol, M., Gašior, J. S., Łaba, J., Korzeniewski, K. and Młyńczak, M. (2023). Evaluation of the performance of gpt-3.5 and gpt-4 on the medical final examination, *medRxiv*. URL: <https://api.semanticscholar.org/CorpusID:259073477>
 - West, C. G. (2023). Advances in apparent conceptual physics reasoning in gpt-4, *ArXiv abs/2303.17012*. URL: <https://api.semanticscholar.org/CorpusID:257913307>
 - Wolfram, S. (2023). What is chatgpt doing . . . and why does it work? <https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/>
 - Yang, Z. and Sinnott, R. O. (2021). Snake detection and classification using deep learning, *Proceedings of the 54th Hawaii International Conference on System Sciences*, Proceedings of the 54th Hawaii International Conference on System Sciences, pp. 1212–1222. <https://scholarspace.manoa.hawaii.edu/server/api/core/bitstreams/f9f2de4d-ddaa-4984-9fc3-9bbca4e27247/content>