AbadIA
Using Reinforcement Learning to play and solve the abbey of crime

Juantomás García - Chief Envisioning Officer - Sngular

s|ngular
it can be done.

Good Morning ING

s|ngular
it can be done.

Who I am

Juantomás García (0-)

- Chief Envisioning Officer @ Sngular
- GDEx2 (Google Developer Expert) for cloud and Machine Learning
- #AbadIA Cheer Leader

Others

- Co-Author of the first Spanish free software book “La Pastilla Roja”
- Former President of Hispalinux (Spanish Linux User Group)
- Organizer of the Machine Learning Spain and GDG Cloud Madrid.

What we will see

How #AbadIA started

My buddy Sebas told me:

Is it possible to create an AI that can learn to play and solve the game?

It’s a really crazy idea with an insane cult 8-bit game, but I like it

The Game:The Abbey of the Crime

THE GAME

Do you know the game?

First 8-bit RPG in pseudo 3D (2.5D)

It was at 1987 and this game is a kind of legend in the video games world.

Based in Umberto Eco book “In the name of the rose”

s|ngular

it can be done.

THE GAME

Do you know the game?

Size of the program: **87 Kb**

That include:

- Code
- Graphics
- Music
- Maps

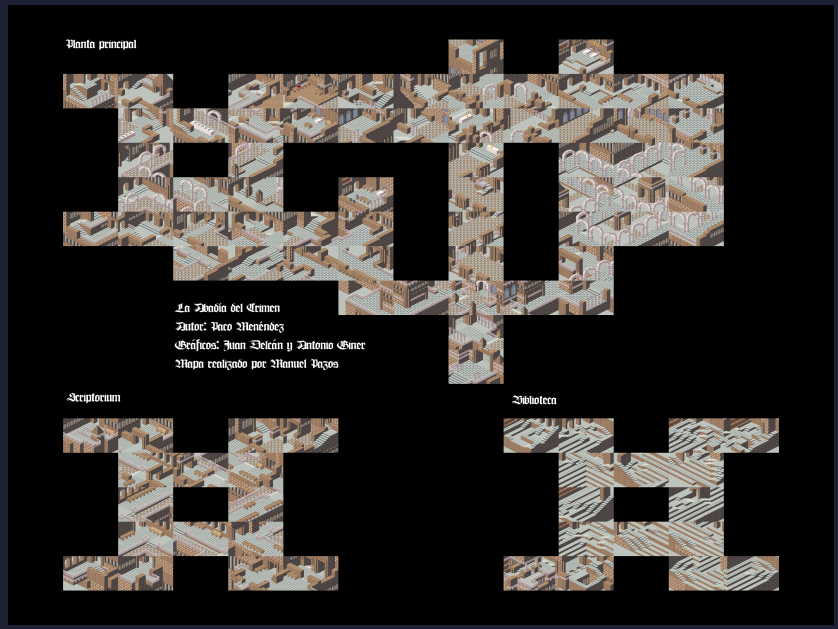

s|ngular

it can be done.

THE GAME

Do you know the game?

The size of this avatar is: 87Kb

s|ngular

it can be done.

ABOUT THE GAME

Was programmed by **Paco Menéndez**

Graphics was designed by **Juan Delcán**

It was programmed using Z80 assembler for Amstrad CPC and Sinclair Spectrum

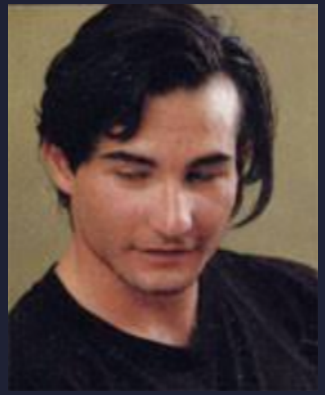

s|ngular

it can be done.

The Plan

The plan:
Make an **AI** that can learn to play and solve the game

What we had

The original executable

Original Code

Manuel Abadía disassembled the original Amstrad CPC game rom.

He documented all the code, marked the graphics, sounds, etc.

GitHub page: https://bit.ly/2I9TMZk

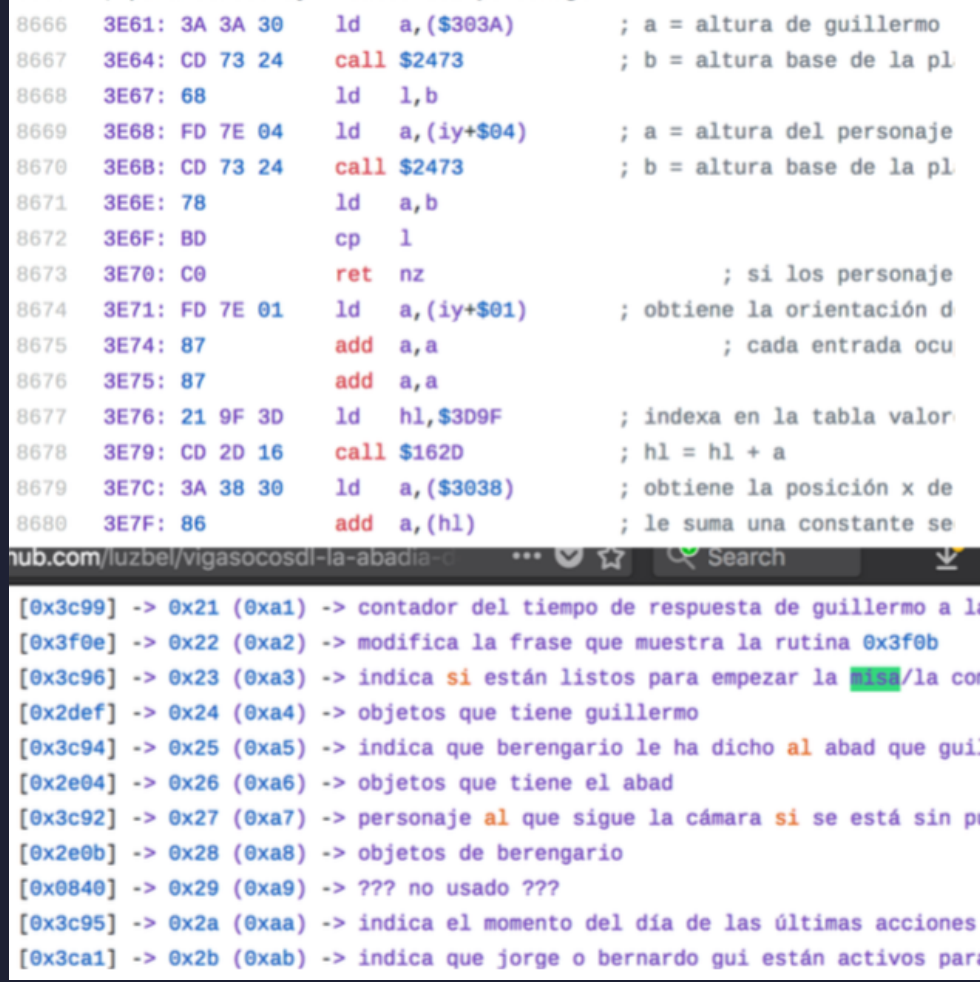

s|ngular
it can be done.

Current Version

GAME EVOLUTION

VigasocoSDL

Manuel Abadía make a port to C++ of the game using his special video games framework: Vigasoco

Sebastian Blanes fork the project to use SDL. VigasocoSDL was born. Now the abbey can be played in chrome, PS3, CC, linux, MacOS, Windows, etc.

GitHub page: https://github.com/luzbel/VigasocoSDL/wiki

How difficult is the Challenge

HOW DIFFICULT IS THE CHALLENGE?

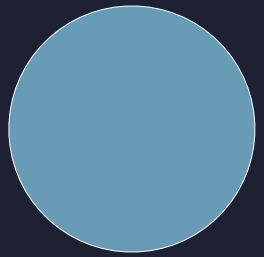

Number of atoms in our universe?

HOW DIFFICULT IS THE CHALLENGE?

Number of atoms in our universe (Hawking said there are more than one universe!!)

HOW DIFFICULT IS THE CHALLENGE?

Number of atoms in our universe (Hawking said there are more than one universe!!)

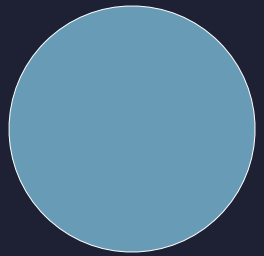

Number of chess legal moves?

s|ngular

it can be done.

HOW DIFFICULT IS THE CHALLENGE?

Number of atoms in our universe (Hawking said there are more than one universe!!)

Number of chess legal moves

HOW DIFFICULT IS THE CHALLENGE?

Number of atoms in our universe (Hawking said there are more than one universe!!)

Number of chess legal moves

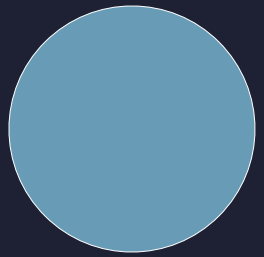

Number of GO legal moves?

HOW DIFFICULT IS THE CHALLENGE?

Number of atoms in our universe (Hawking said there are more than one universe!!)

Number of chess legal moves

Number of GO legal moves

s|ngular

it can be done.

HOW DIFFICULT IS THE CHALLENGE?

Number of atoms in our universe (Hawking said there are more than one universe!!)

Number of chess legal moves

Number of GO legal moves

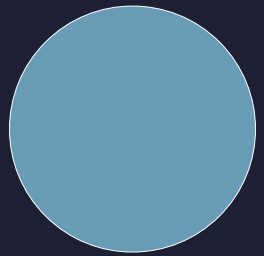

AbadIA moves needs to solve the game?

HOW DIFFICULT IS THE CHALLENGE?

Number of atoms in our universe (Hawking said there are more than one universe!!)

Number of chess legal moves

Number of GO legal moves

AbadIA legal moves where N is the depth of the game you wish to check. For N= 10000 is 10^10000

s|ngular

it can be done.

Just for the guys that not attended enought math classes.

Universe

100000000000000000000000000000000000000000000000000000000000000000000000000000000

Chess

1000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000

GO

100000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000
000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000
000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000
000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000
000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000000
00000000000000000000000000000000000000000000000

The Abbey of Crime

1000
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00
00000000000000000000000000000000000000000000000000

singular

it can be done.

The AI & Human Players only know if the game is solved after 10.000 actions.

HOW DIFFICULT IS THE CHALLENGE?

s|ngular

it can be done.

HOW DIFFICULT IS THE CHALLENGE?

s|ngular
it can be done.

HOW DIFFICULT IS THE CHALLENGE?

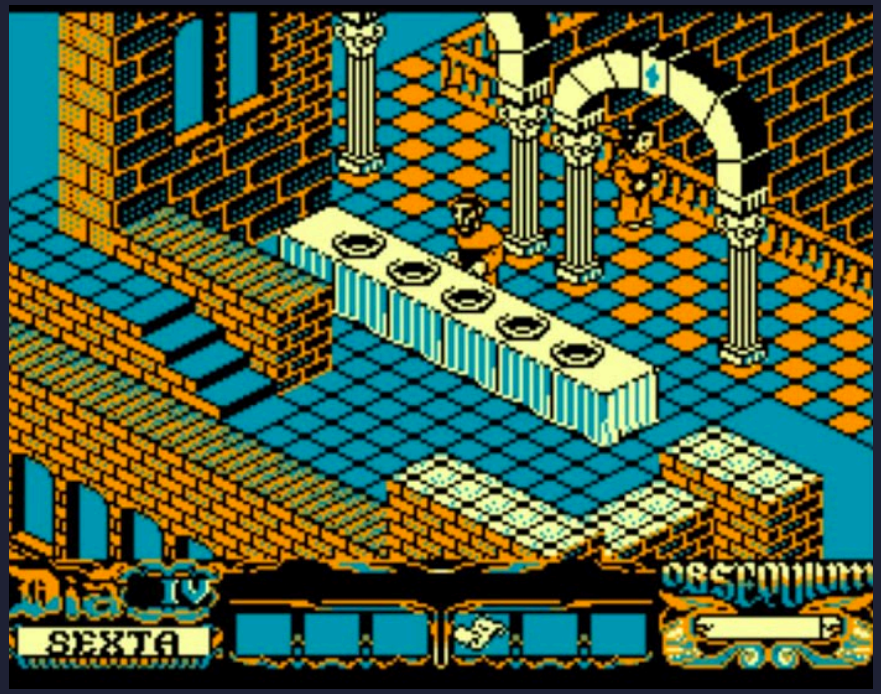

s|ngular
it can be done.

HOW DIFFICULT IS THE CHALLENGE?

s|ngular
it can be done.

First Steps

Interacting with the game

AI TOOLS

You will need lots of tools before starting to create the AI

1. How to interact with the game.
2. How you get and save the game information.
3. How to play millions of games.

s|ngular
it can be done.

HOW TO INTERACT

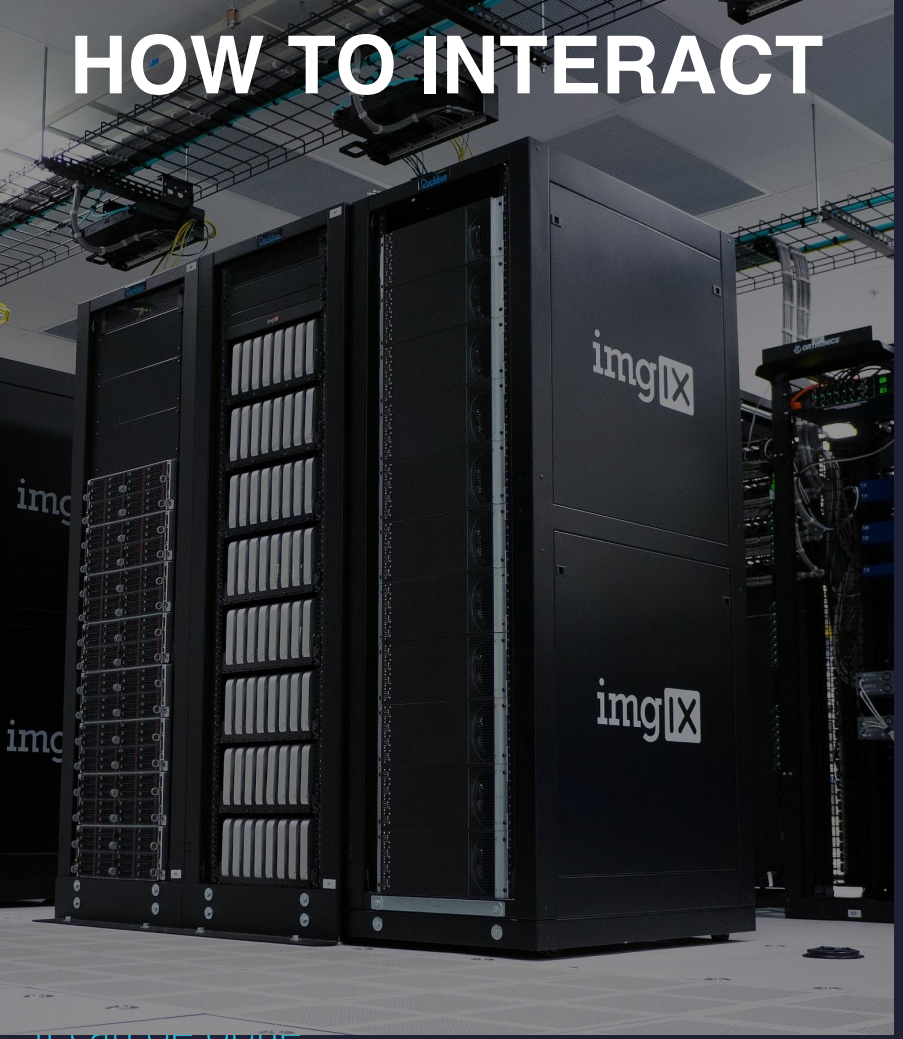

You will need send actions, get game info, reset it, etc.

How to interact with the game??

HOW TO INTERACT

First version with too Beers

HOW TO INTERACT

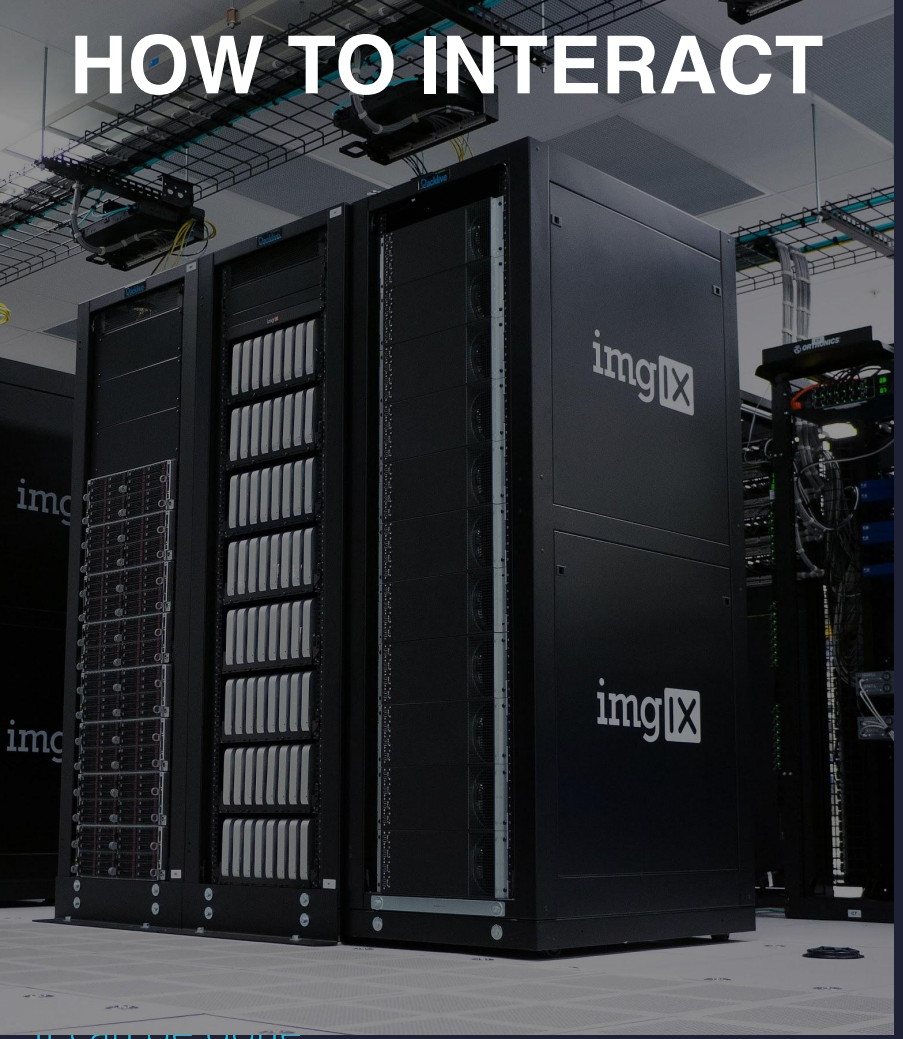

You may use a fake keyboard driver, a unix pipe, a socket, but

do you know what is the really cool way to do this:

HOW TO INTERACT

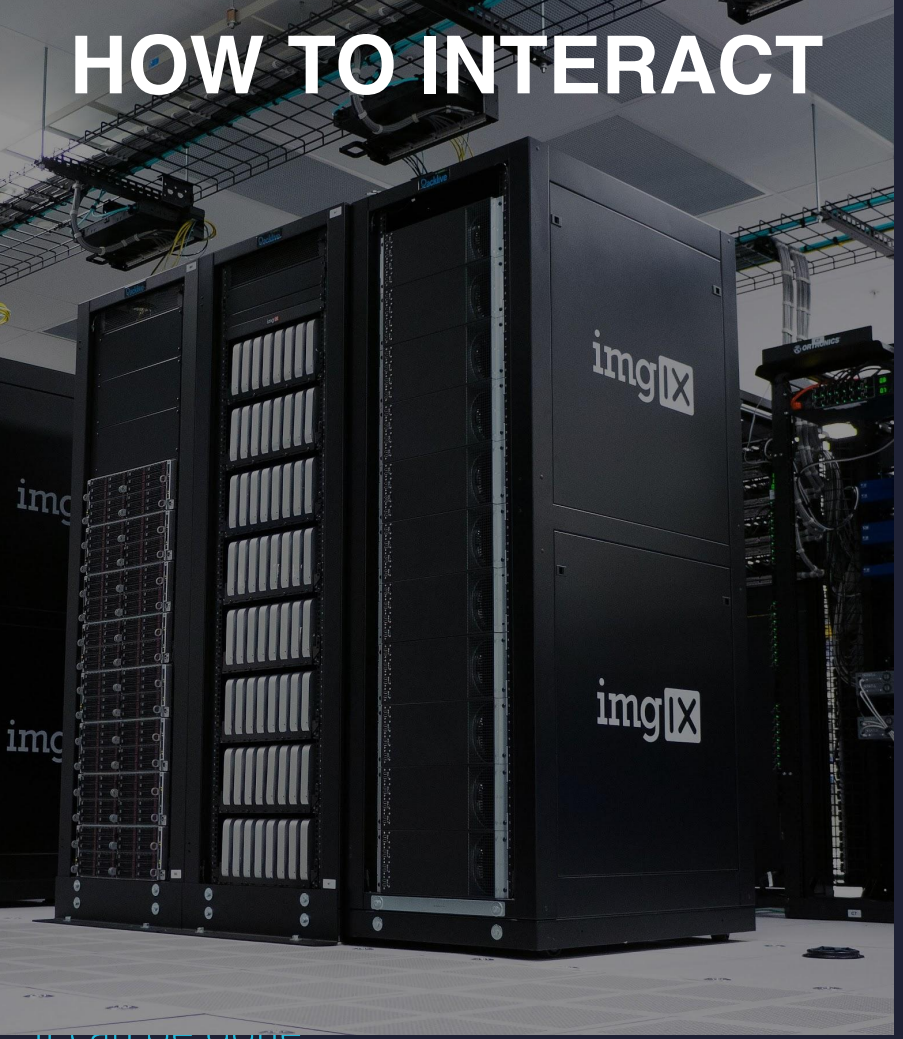

An embedded web server inside the game !!!!

HOW TO INTERACT WITH THE GAME

Now we have a REST based API.
Two way communication.

We send actions
- Moves
- Resets
- Save/Load states

We got information
- State dumps
- Actions
- Checkpoints

s|ngular
it can be done.

AS EASY AS

```
curl -X POST http://localhost:4477/abadIA/game/current/actions/RIGHT
```

or

```
curl -X GET http://localhost:4477/abadIA/game/current –H 'accept:
application/json’
```

or

Use python **request** library

Demo game engine + web server + curl

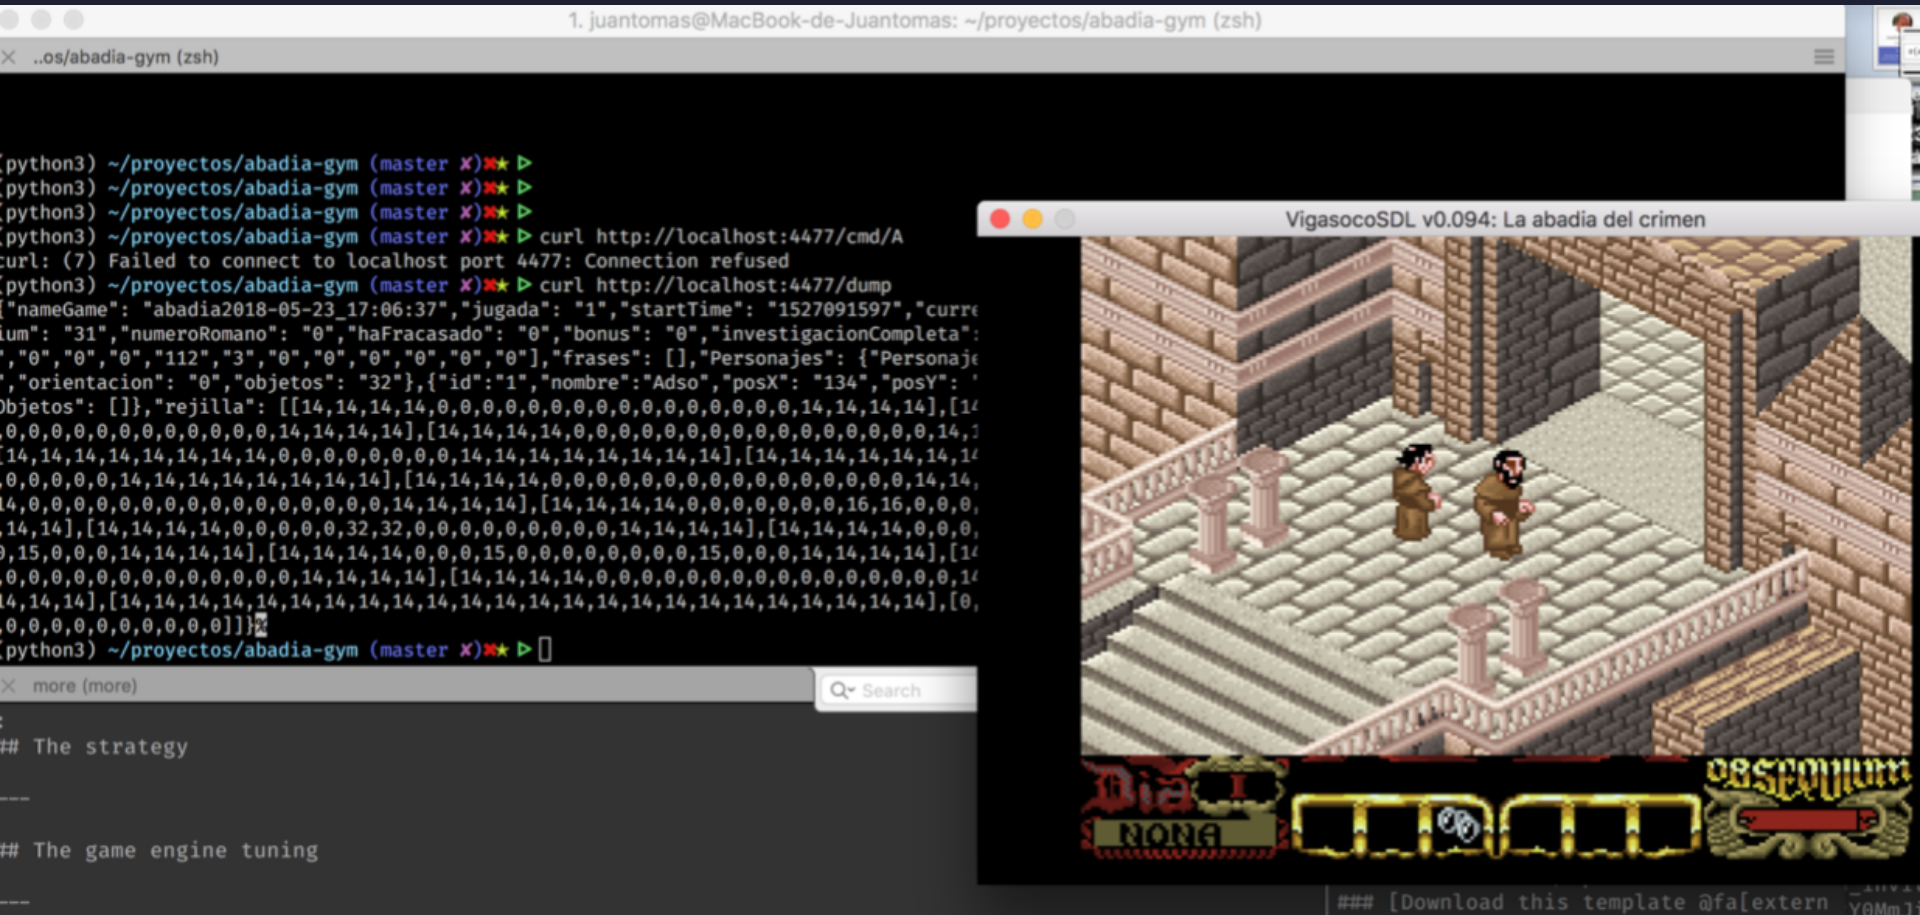

Gathering Information

HOW TO SCALE IT

At the beginning a laptop was enough.
But very soon you need more CPU/ GPU.
Then product like Google Cloud is your best ally.
We had created a few Dockers, so now we can execute lots of instances of the game in parallel.
note: If we use Google Cloud services like GKE, we can launch hundred of games in parallel.

s|ngular
it can be done.

GATHERING INFORMATION

We recollect a lot of information

- Game Info (timestamps, rewards, bonus, obsequium)
- Games moves (state, action, reward, new_state)
- Checkpoints (to restore the game at an interesting time)
- ML Models (for recovering good models o just make a benchmark)

```
(python3) ~/proyectos/abadia-gym (master ✗)✖★ ▷ curl http://localhost:4477/dump
{"nameGame": "abadia2018-05-23_17:06:37","jugada": "1","startTime": "1527091597","currentGame": "1527091626","
ium": "31","numeroRomano": "0","haFracasado": "0","bonus": "0","investigacionCompleta": "0","porcentaje": "0",
","0","0","0","112","3","0","0","0","0","0","0"],"frases": [],"Personajes": {"Personaje": [{"id":"0","nombre":
","orientacion": "0","objetos": "32"},{"id":"1","nombre":"Adso","posX": "134","posY": "169","altura": "0","ori
Objetos": []},"rejilla": [[14,14,14,14,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,14,14,14,14],[14,14,14,14,0,0,0,0,0,0,0
,0,0,0,0,0,0,0,0,0,0,0,0,0,0,14,14,14,14],[14,14,14,14,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,14,14,14,14],[14,14,14,14,1
[14,14,14,14,14,14,14,14,14,0,0,0,0,0,0,0,0,14,14,14,14,14,14,14,14],[14,14,14,14,14,14,14,14,14,0,0,0,0,0,0,0,0,14,
,0,0,0,0,0,14,14,14,14,14,14,14,14],[14,14,14,14,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,14,14,14,14],[14,14,14,14,0,0
14,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,14,14,14,14],[14,14,14,14,0,0,0,0,0,0,0,0,16,16,0,0,0,0,0,0,0,14,14,14,14],[1
,14,14],[14,14,14,14,0,0,0,0,0,32,32,0,0,0,0,0,0,0,0,0,14,14,14,14],[14,14,14,14,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0
```

GATHERING INFORMATION

It takes a lot of time to get all the parts working all together.

Building tools, testing every piece, every option.

Sometimes I feel like I was Mario Bros.

s|ngular
it can be done.

Playground for the game

PLAYGROUND FOR PLAYING

One the most frequently used tool is OpenAI Gym

➡

So we design an AbadIA gym

➡

The gym is a standard place to train and interact with Reinforcement Learning agents.

➡

In our project the gym is framework to wrap the game engine.

s|ngular
it can be done.

Github Repo: https://www.github.com/LaAbadIAdelCrimen/ abadia-gym

SO WE HAVE

So what's the next step

An small intro to Reinforcement Learning

s|ngular
it can be done.

CREATE A RL AGENT

A RL agent is a program that interacts with an environment, in our case a OpenAI gym for AbadIA, and learn from observations and rewards.

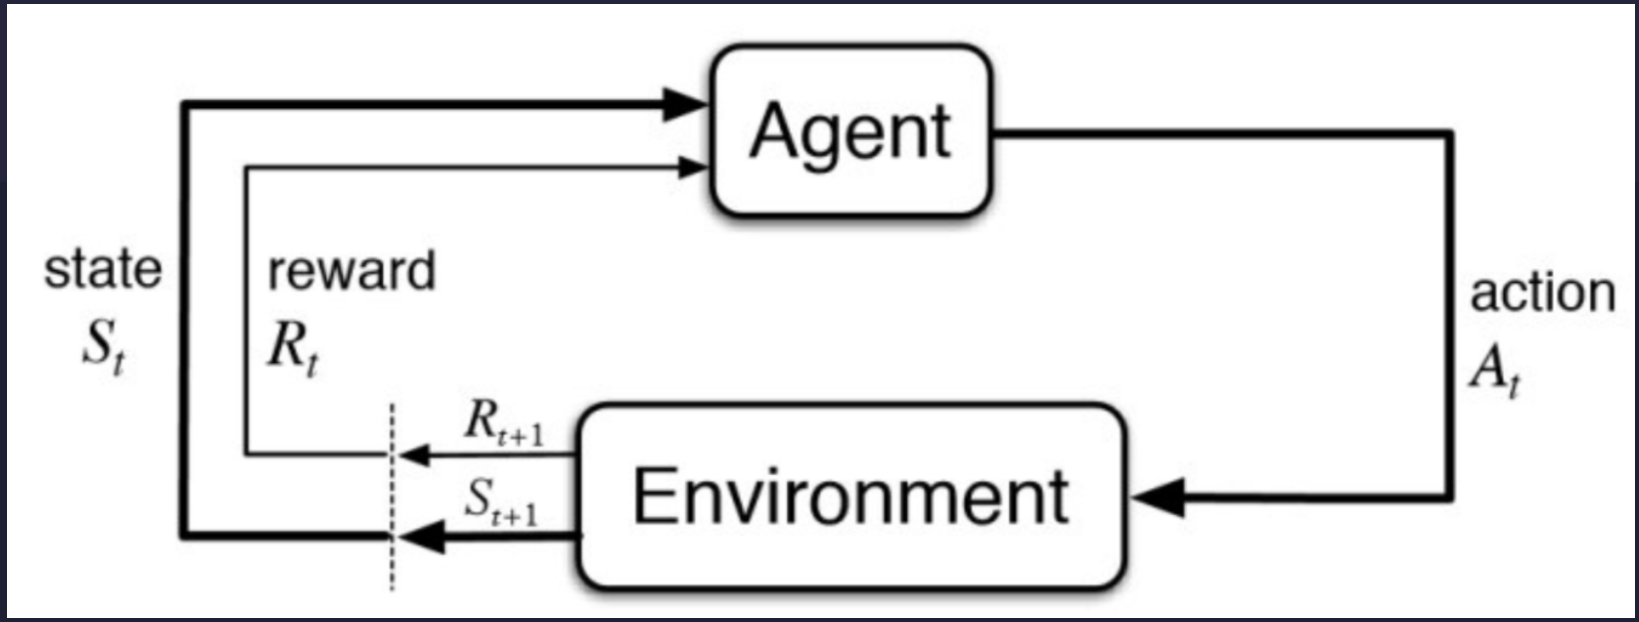

And that looks like:

Demo Time + Video

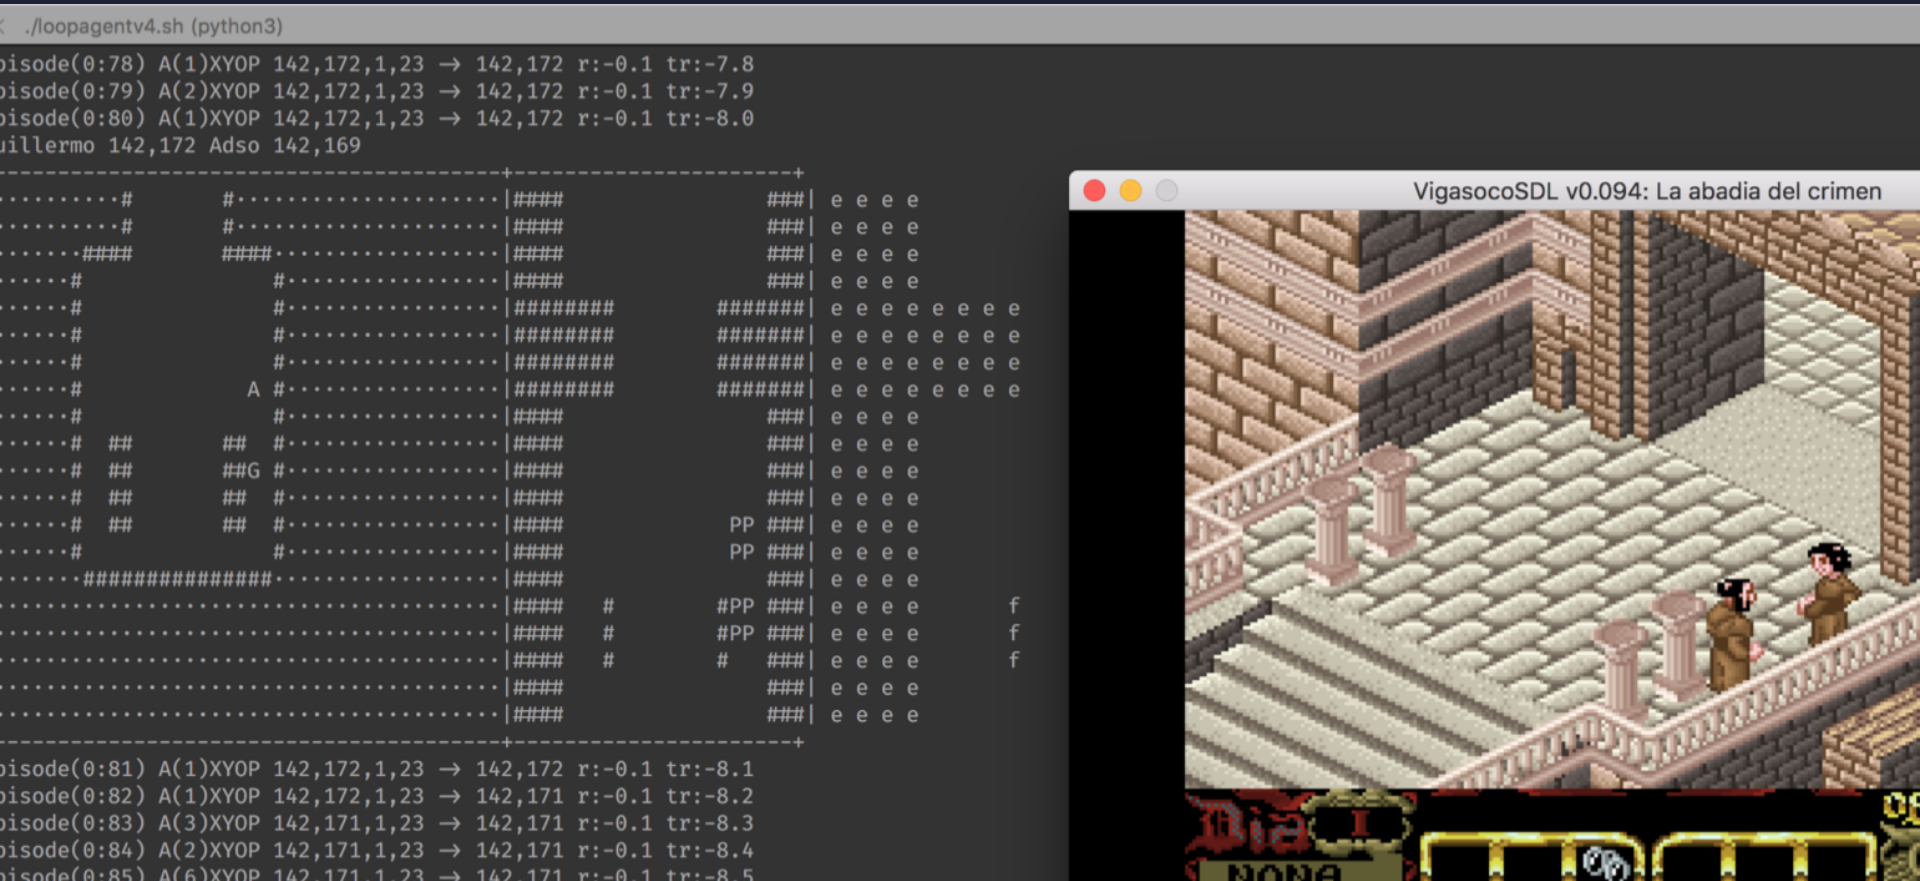

AbadIA RL Strategies

s|ngular
it can be done.

REINFORCEMENT LEARNING STRATEGIES

- Random
- Q*tables
- Q*Learning
- DQN
- CNN
- A3C
- MiniGames

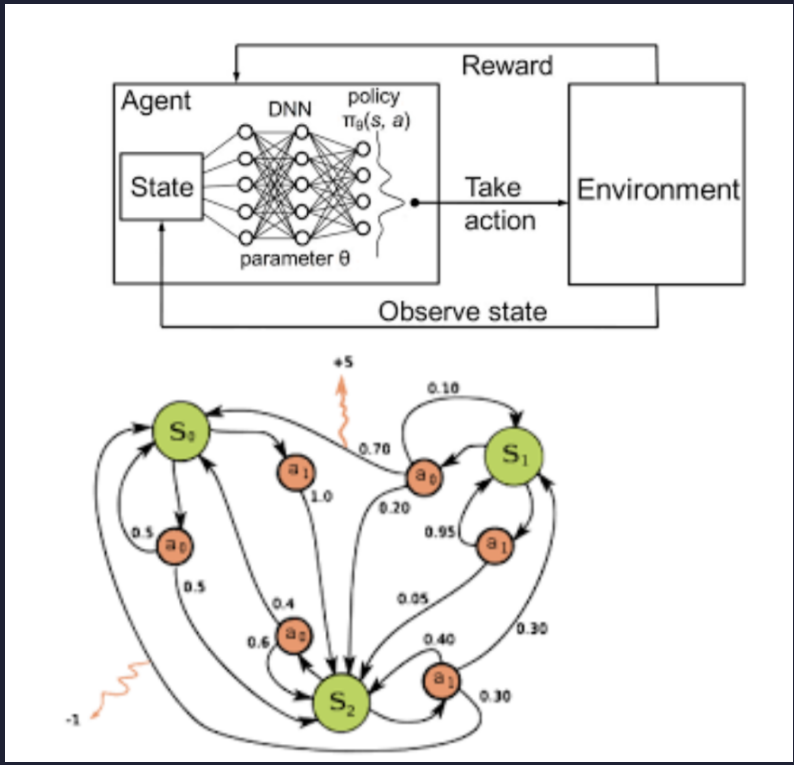

s|ngular

it can be done.

Simple Neural Network Model

```
def create_model(self):
    self.env.logging.info("Creating a new model v4")
    model = Sequential()
    state_shape = self.env.observation_space.shape
    model.add(Dense(24, input_dim=state_shape[0], activation="relu"))
    model.add(Dense(48, activation="relu"))
    model.add(Dense(24, activation="relu"))
    model.add(Dense(self.env.action_space.n))
    model.compile(loss="mean_squared_error",
        optimizer=Adam(lr=self.learning_rate))
    return model

def load_model(self, name):
    self.env.logging.info("Loading a model from: ({})".format(name))
    return load_model(name)
```

AbadIA
Architecture Parts

s|ngular
it can be done.

Game Engine & Agent Containers

s|ngular
it can be done.

First we create the first 2 containers for:

- AbadIA Game Engine
- AbadIA Agent

s|ngular
it can be done.

K8S small cluster

Cluster Features:

- Preemtible
- n1-standard-1 (7€ month)

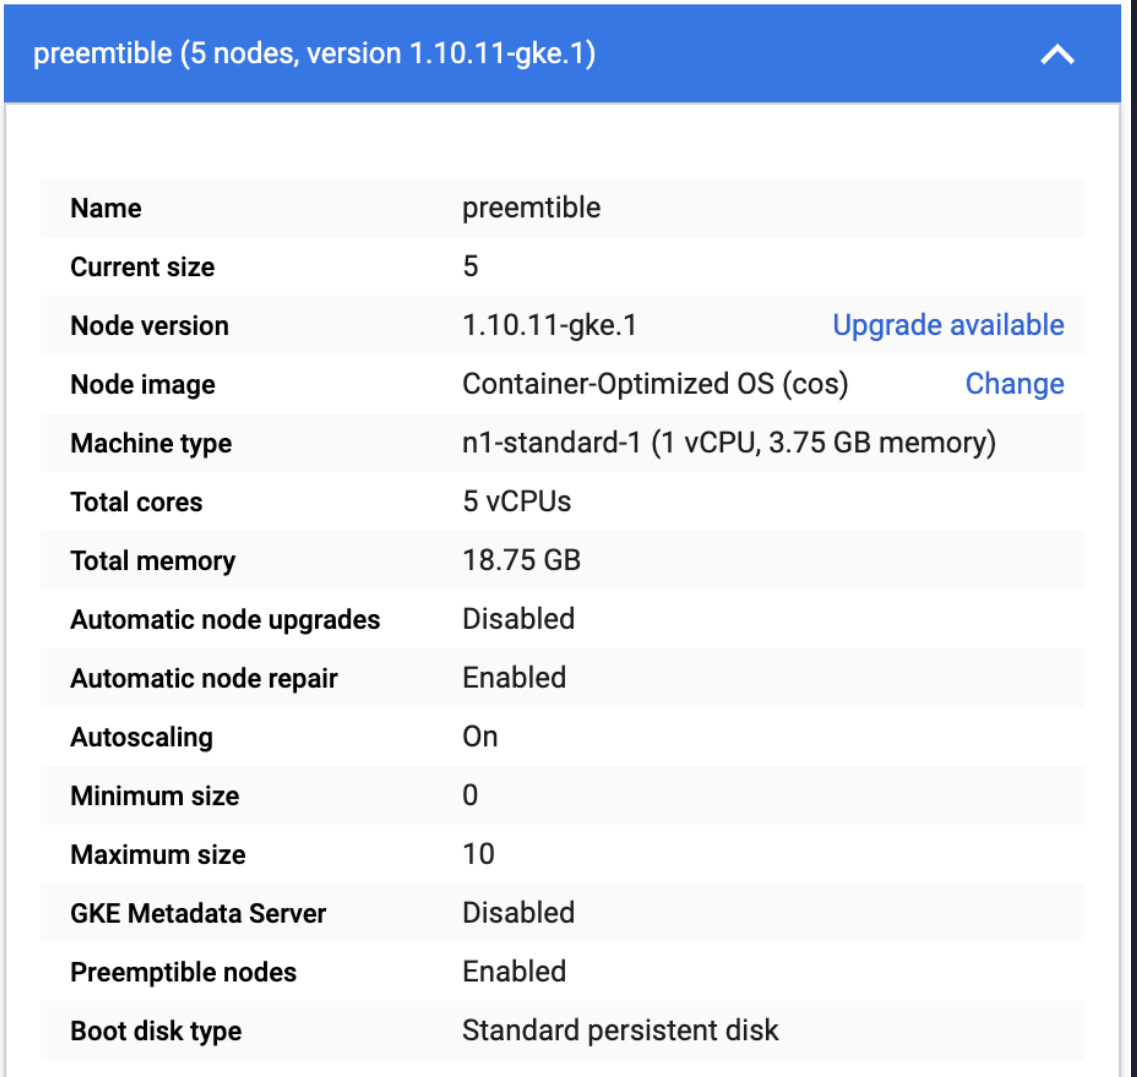

s|ngular
it can be done.

Gathering Information

Game info pipeline

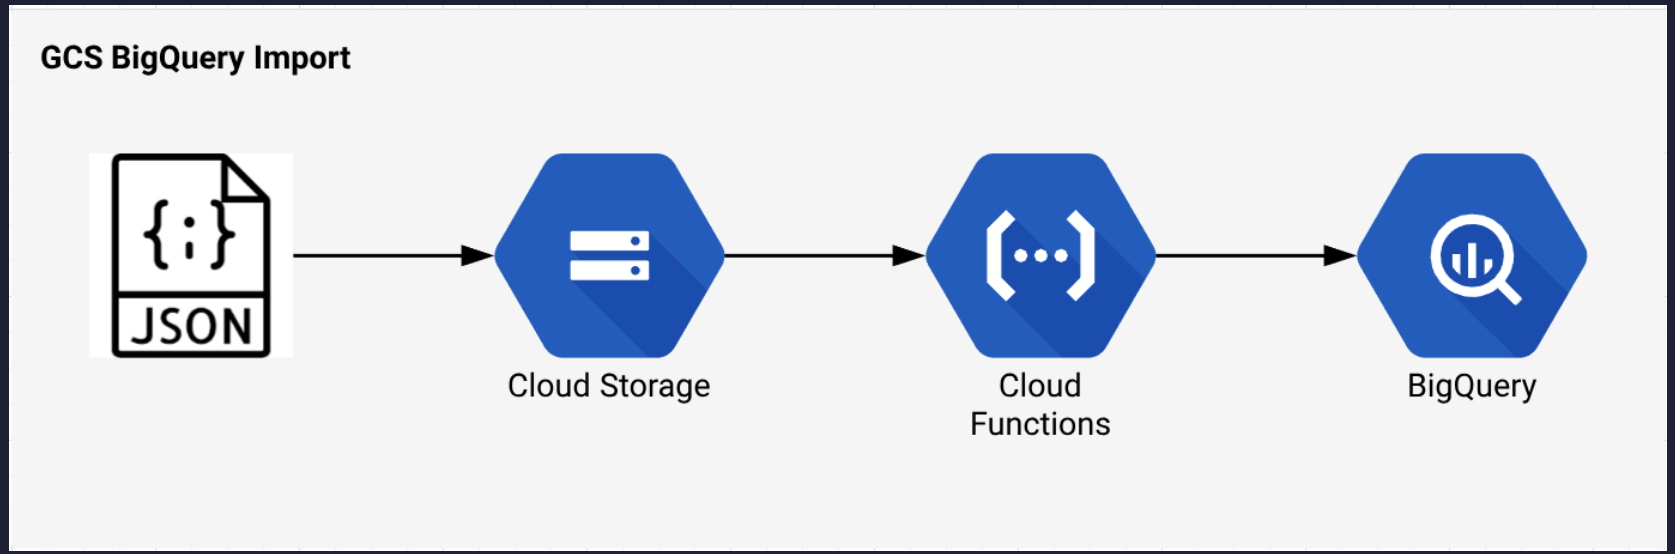

AutoGaming

s|ngular
it can be done.

Games Checkpoints

Game Generator

Cloud Pub/Sub

s|ngular
it can be done.

Storage of the info:

- Game Information
- Actions - States - Rewards
- Snapshots
- Models

s|ngular
it can be done.

How to Integrate the all the information:

- Serverless
- Easy
- Fast

s|ngular
it can be done.

Advanced and Massive Analytics:

- Game Information
- Actions - States - Rewards
- Snapshots
- Even some ML (LRs)

Google BigQuery

s|ngular
it can be done.

Notebooks

How to manipulate and visualize with Datalabs or Colab:

- Notebooks
- Managed
- Integrated with Google Drive / Bigquery

s|ngular

it can be done.

AI Frameworks

How to training and serve models at scale:

- Tensorflow / Keras
- ML Engine
- AutoML

s|ngular
it can be done.

Orchestrating

How to execute
hundreds of instances:

- Kubernetes
- Google Containers

s|ngular
it can be done.

Show me the code (how to collaborate)

HOW TO COLLABORATE

GitHub
LaAbadIAdelCrimen

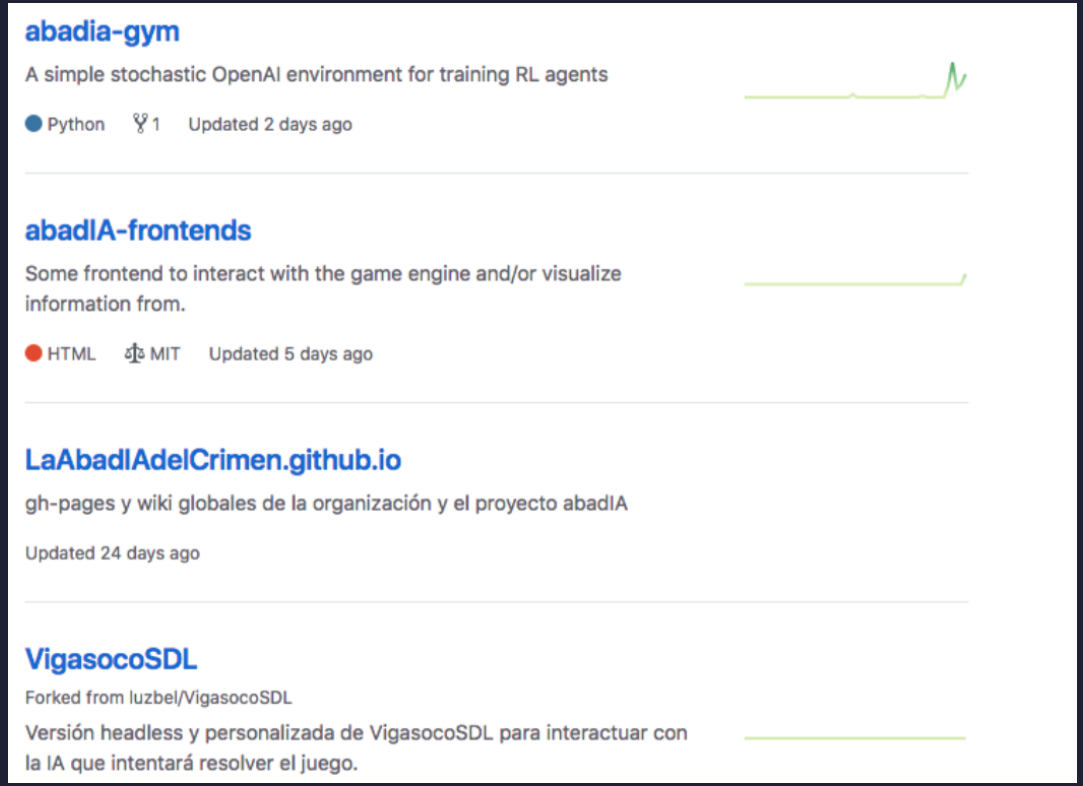

s|ngular
it can be done.

https://github.com/LaAbadIAdelCrimen

Join US, we'll share the glory with you

Slack AbadIA

JOIN US!!!!

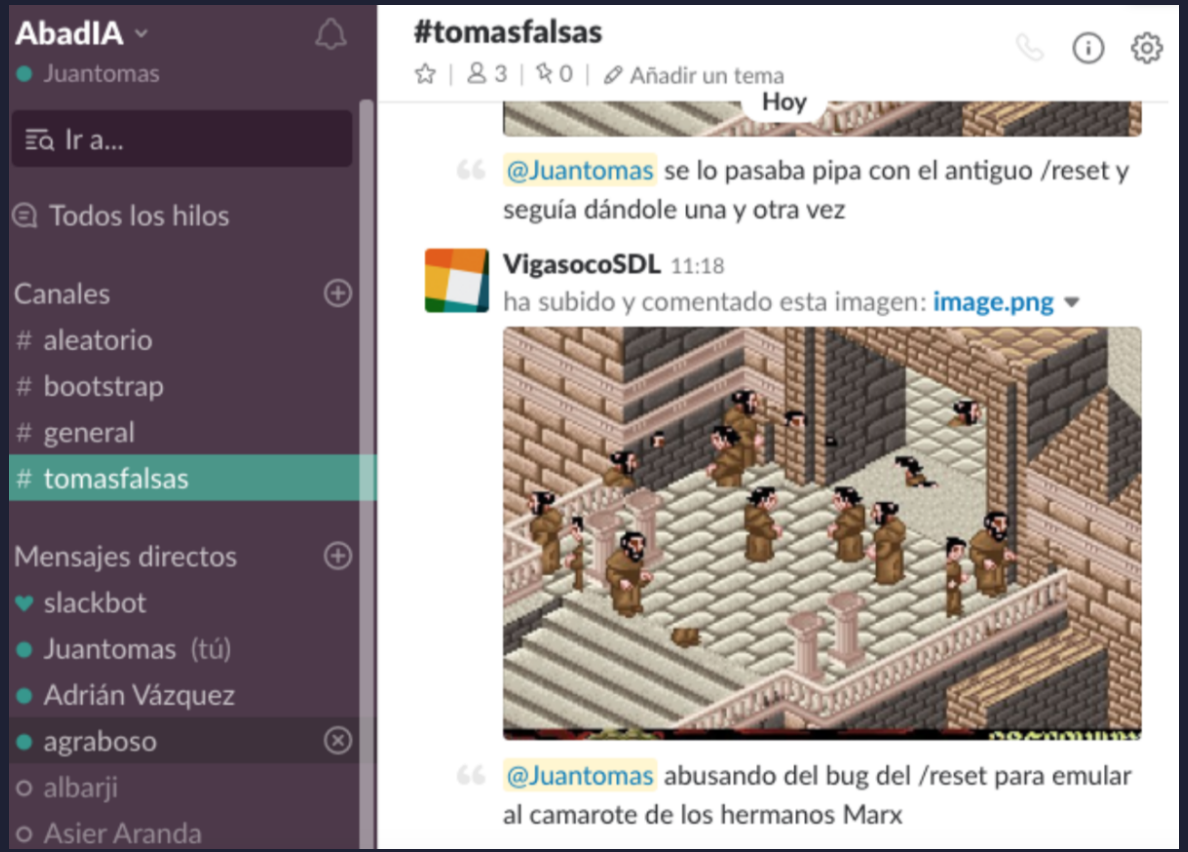

s|ngular

it can be done.

Questions?

twitter: @juantomas
juantomas.garcia@sngular.com

We're Hiring, Sngular People

Selfie Time: If you like the talk just smile while I take the selfie ;-)

This talk have a free questions lifetime warranty: If you have any questions or concerns about this talk, feel free to contact me anytime.

s|ngular
it can be done.

Thank You!!

s|ngular
it can be done.