

Antigravity: the assistant that works on your computer

Projects, permissions, models and review: your first real job, step by step

Gianluca Bonomo · www.gianlu.it · info@gianlu.it



Contents

How to use this handout	3
1. From chat to agent.....	3
2. The model and the tank.....	4
3. Projects and environment: where the agent works.....	6
4. The first real job, step by step.....	9
4.1 Step 1: the permission request.....	10
4.2 Step 2: the work.....	11
4.3 Step 3: the result.....	12
5. Reviewing the work: the Review panel.....	13
6. Commands with the slash and the at sign	14
7. Scheduled tasks	16
8. Safety: the settings that matter	17
9. What it is really for	19
Hands-on exercise	20
Questions that come up in class.....	21
Glossary.....	21
Further reading	22

How to use this handout

This handout gathers, in written and ordered form, everything covered in the lesson. You do not need to memorise it: it is there so you can find things again calmly, later, when the lesson is over and a question is left hanging.

GOAL

Understand what Antigravity is, create a project on a dedicated folder, assign a task, manage permissions and review the result safely.

What you will be able to do


- Explain the difference between a chat and an agent that works on the computer.
- Create a project on a dedicated folder and choose the model.
- Read and answer the agent's permission requests.
- Review changes with the Review panel and check settings and limits.

1. From chat to agent

A chat gives you advice. An agent gives you the job done.

So far Gemini has written texts, created images, answered questions. But if it tells you «create a web page with the pond checklist», you still had to create the file yourself.

Antigravity takes one step further: it is an **agent**, that is an assistant that can **act on your computer**. It reads the files in a folder, creates new ones, runs commands and reports what it did. Think of it as a **site foreman**: you say what you want, it organises and works, and asks permission before anything delicate.



Antigravity's start screen. Previous conversations in the sidebar have been blurred.

According to Google, **Antigravity 2.0** is a standalone desktop application, available for macOS, Linux and Windows, built to work with agents and without the components of a code editor. Besides the app there are other pieces of the same world:

Piece	What it is
Antigravity 2.0 app	The window you see here: conversations, projects and agents.
Antigravity IDE	A separate code editor, for those who want to write programs; from the app it is opened with the Install IDE button at the top right.
Antigravity CLI	The command-line version, for those who use the terminal.
Antigravity SDK	The tools for developers, to build their own agents.


The pieces of the Antigravity family according to Google (May 2026). In this course we use only the app.

IN PLAIN WORDS

A chat answers. An agent works. That is why permissions and a bit more caution are needed.

2. The model and the tank

Under the writing box there is the name of the model, for example **Gemini 3.8 Flash High**. Click it (or press **Ctrl+ /**) and the list opens.




Antigravity's list of models, on 7 October 2026.

On my account, that day, there were these models. Note that they are not only Google models.

Model	Note in the list
Gemini 3.8 Flash (High)	The one selected by default.
Gemini 3.7 Flash and 3.6 Flash (Medium)	Marked «Leaving Soon»: they are about to go.
Gemini 3.1 Pro (Low)	The Pro model of the Gemini family, with reasoning set to Low.
Claude Sonnet 4.6 and Claude Opus 4.6 (Thinking)	Anthropic models, with the label «Notice».
GPT-OSS 120B (Medium)	An open model by OpenAI, with the label «Notice».

The models visible in the selector. The list changes often.

The second word (High, Medium, Low) shows **how hard the model tries to reason**. The higher it is, the more accurate and the more it uses. For this course's tests I kept Flash High.




Settings, Models & Usage: the plan, the credits and the usage limits.

- **Your Plan:** my plan was Google AI Pro, with the Upgrade button to move to Ultra.
- **Gemini Models:** two limits, **Five Hour** (short) and **Weekly**, with the percentage left and when it refills. That day I had 98 per cent of the week.
- **Claude and GPT models:** they have separate limits.
- **Model Credits:** the *Enable AI Credit Overages* switch is used to spend paid credits when the quota runs out. It was off, and for a beginner it is better to leave it so.

3. Projects and environment: where the agent works


Before writing the task you must say **where** the agent can work. Above the box there is the **No Project** menu.



The projects menu: New Project, Quick Start, No Project.

- **No Project:** no working folder. The agent works in a more limited way.
- **New Project:** you choose a folder on the computer, which becomes the project.
- **Quick Start:** a quick start to begin without configuring.

I chose **New Project** and pointed to a **new, empty folder**, created for the purpose and called **demo-antigravity**. The normal Windows window for choosing a folder opens.




The selected project: the folder name appears at the top, the environment menu under the box.

THE DEDICATED FOLDER RULE

An agent can read and write in the project folder. **Never point to the Documents folder or the whole disk:** create a folder for each job, ideally empty or with a copy of the files. It is the fence the agent can move within.

Under the box there is also the **Local** menu (**Ctrl+.**). It is used to choose the **environment**.




The environment: Local or New Worktree (only for Git folders).

Local works directly in the folder. **New Worktree** creates a separate working copy, but is available **only for Git repositories** (the file history of programmers). For normal jobs Local is fine.

4. The first real job, step by step

The work cycle




The work cycle with an agent

Let's do a simple, recognisable task. I wrote, in Italian:

Create an index.html page with a checklist to build a 3-by-2-metre garden pond, with boxes to tick and a simple style.

The prompt given to Antigravity (translated)




The request written in the box, with the project selected.

I press Enter. Antigravity gives the conversation a title (here «Creazione Checklist Laghetto Giardino», Creating a Garden Pond Checklist) and starts.

4.1 Step 1: the permission request

After a few seconds the agent wants to look at what is in the folder. Before doing so it **stops and asks**.




The permission request: «Allow listing directory contents?» with the command it wants to run.

Option	What it means	When to choose it
1. Yes, allow this time	Yes, just this once.	Almost always, at the beginning.
2. Yes, and always allow ... in this conversation	Yes, and next times too, in this conversation.	For harmless commands that repeat.
3. ... in this project	Yes, always in this project.	Only if you trust the project.
4. Yes, and always allow ...	Yes, always, everywhere.	With great caution, never for commands that delete.
5. No (tell the agent what to do instead)	No, and you tell the agent what to do instead.	When the command does not convince you.

The five answers to the permission request.


The command here was **Get-ChildItem -Path .**, that is «list the files of this folder»: harmless. I chose **1**.

4.2 Step 2: the work



The agent reasons («Thinking for 12s») and shows the steps it is carrying out.


Antigravity shows what it is doing: «Ran Get-ChildItem», «Thinking...». You can open its reasoning. After about a minute the file was created: **index.html**, about 20 kilobytes, with almost six hundred lines.



The agent wrote the file («Edited index.html +592 -0») and rereads it to check; then it asks a second permission.

Note something I liked: the agent **rereads its own work** in pieces («Analyzed index.html #L1-100») and runs a check with **Test-Path**, asking permission again. I answered **1** again.

4.3 Step 3: the result



The agent's final summary, with the Review button and «1 file changed +593 -0».


At the end the agent writes a **report**: what it did (checklist divided into six phases, progress bar, automatic saving of ticked boxes, quick buttons, green responsive style) and how many changes it made: **1 file changed, +593 -0**, that is one file with 593 lines added and none removed.



The page created, opened in the browser: title, progress bar and the first items with boxes to tick. Text in Italian.

5. Reviewing the work: the Review panel

The **Review** button opens on the right a panel with the **changes line by line**: in green what was added, in red what was removed. It is your safety net.



The Review panel: the code of the index.html file with line numbers, all green because it is a new file.

1. Read the agent's report: does it say what it did?

Antigravity: the assistant that works on your computer


2. **Open the result** (here, the page in the browser) and really try it.
3. **Check numbers and facts.** In my page, for example, the EPDM liner calculation comes to about 5.6 metres: redoing the sum with the depth you chose yourself, the result changes. An AI writes well, but sums must be redone.
4. **Ask for corrections** in the same conversation: «Add a section on costs», «Change the green to blue».

THE IDEA IN ONE LINE

The agent works, you check. A good foreman never signs off a job without looking at it.

6. Commands with the slash and the at sign

In the writing box you can type **/** to see the **commands**, or **@** to call something. They are shortcuts to tell the agent **how** to work.




Typing **/** opens the list of commands.

Command	Description (as shown by Antigravity)	What it is for
/goal	Run until the specified goal is completely finished.	You give a goal and the agent works until it is done.
/schedule	Run an instruction on a recurring schedule or as a one-time timer.	Schedule a recurring task or one with a timer.

Command	Description (as shown by Antigravity)	What it is for
/browser	Invoke a browser agent for web tasks.	An agent that uses the browser for web tasks.
/plan	Plan carefully before executing a task.	First a plan, then the work: ideal for big tasks.
/grill-me	Interview me to align on a plan.	The agent asks you questions before starting.
/teamwork-preview	Invoke a team of agents to autonomously tackle large projects.	A team of agents for big projects.
/learn	Reflect on recent successes or corrections to capture reusable skills or rules.	Learns from your approvals and corrections.
/boost	Invoke the Boost multi-agent orchestrator for complex tasks.	A coordinator of several agents for complex tasks.


The first commands of the list on 7 October 2026; the list continues when scrolling.

The beginner's advice: for any task that is not trivial, start with **/plan** or **/grill-me**. You spend two minutes getting aligned and save half an hour of corrections.



Typing @ lets you call the Rules or a previous Conversation.

With @ you call the **Rules**, your work rules, or a **Conversation**, that is a previous conversation to resume from. At the bottom right of the box there is also the **Main Agent** menu.




The Agents menu: the main agent and, on my computer, two specialised agents.

Main Agent is the basic agent. Other specialised agents may appear in the list (on my PC, two I already had installed): each is a collaborator with a precise skill, which the main agent can call.

7. Scheduled tasks

As in Gemini, in Antigravity too there are tasks that start by themselves. You find them in the sidebar with **Scheduled Tasks** (shortcut **Ctrl+U**).




The Scheduled Tasks page: search, New button and, here, no task configured.

On my computer, that day, it said «No scheduled tasks configured». With **New** or with the **/schedule** command you create one, for example «every Monday morning check whether there are new checklists in the folder and summarise them». These tasks use up limits: **few and precise**.

8. Safety: the settings that matter

At the bottom left there is **Settings**. The first page, **General**, contains the rules of the game.




Settings, General: execution, global permissions, agent behaviour and network permissions. Account name and email are blurred.

Setting	Value on my PC	What it governs
Queued Messages	Queue	Whether messages sent while the agent works are queued (Queue) or sent immediately (Send Immediately).
Security Preset	Default	The general level of actions the agent can take.
Tool Permissions	7 rules	Permissions for file, terminal and MCP tools.
Plan Review Policy	Always Ask	The agent always asks you to review its documents (plans).
Network Access Rules	Open	The web addresses it may or may not read.

The main settings of the General page.

Another useful page is **Customizations**: it shows your **skills**, **subagents** and **MCP** tools, with a **token budget** (here 20,000). If you exceed it, Antigravity automatically excludes the larger customisations.



Customizations: the token budget and the list of skills. A warning says the budget has been exceeded.

THE CAUTION CHECKLIST

- 1) A dedicated folder for each project.
- 2) Security preset **Default** and «Always Ask» until you really know what you are doing.
- 3) For permissions, **«allow this time»**, not «always».
- 4) Never personal or confidential data in the folder.
- 5) A backup copy of important files before letting the agent work.
- 6) Read the Review panel before trusting.

9. What it is really for

1. **Pages and small tools:** an interactive checklist, a calculator, a page for your shop.
2. **Tidying and transforming files:** reorganising a folder of documents, merging tables, renaming files in bulk.
3. **Checking and summarising:** reading a folder of documents and preparing a summary.
4. **Learning to program:** asking the agent to explain line by line what it wrote.

It is not the right tool for everything: to write a text or a piece of research, chat or Canvas are enough and use less. Antigravity is for when there is **a job to do on the computer**.

Hands-on exercise

Thirty minutes. You need the Antigravity app installed and an empty folder created for the purpose.

1. Create on your PC a new, empty folder, for example *antigravity-test*, away from your documents.
2. Open Antigravity, click No Project, choose New Project and point to that folder.
3. Write a simple task, for example: «Create an index.html page with a checklist for my hobby, with boxes to tick».
4. When the agent asks for a permission, read the command and answer 1, Yes allow this time.
5. Open Review, look at the added lines and then open index.html in the browser.
6. Ask for a change in the same conversation (a colour, an extra section) and check again.
7. Open Settings, then Models, and look at your plan and limits: how much have you used?

WATCH OUT

Did you choose «allow this time» for every permission?

Did you open and try the result before trusting it?

Questions that come up in class

Do I need to know how to program to use Antigravity?

Not for simple jobs: describe what you want in plain language. But you must be able to check the result, and that takes a bit of curiosity and a lot of caution.

What is the difference between the app and the IDE?

The Antigravity 2.0 app is built to work with agents, without a code editor. The IDE is a separate code editor, which you install from the Install IDE button.

Can I use models other than Gemini?

In my list there were also Anthropic models (Claude) and an open OpenAI model. The list depends on plan and period.

Can the agent delete my files?

If you give it permission, in theory yes. That is why you work in a dedicated folder, with backups, answering «allow this time» and reading every command.

Is it free?

It works within the limits of your Google AI plan (Five Hour and Weekly) and, if you turn it on, with paid credits. Check the Models & Usage page.

Glossary

Term	What it means
Agent	An AI assistant that does not just answer: it acts, for example by reading and writing files on your computer.
MCP	A standard way to connect external tools and services to an agent.
Permission	The confirmation the agent asks you for before running a command.
Project	The computer folder in which the agent may work.
Review	The panel that shows, line by line, what the agent changed.
Subagent	A specialised agent that the main agent can call for a precise task.

Further reading

Google, Introducing Google Antigravity 2.0 - May 2026 announcement (desktop app, IDE, CLI, SDK)

Antigravity app on the author's PC: models, settings, test project - screenshots and tests of 7 October 2026

WATCH OUT

Data checked on 7 October 2026 on the installed Antigravity app. The model list and menu entries change with updates. The job shown was done in a dedicated, empty folder created for the purpose. The sums in the sample page were not verified and must be redone.