SIG & Territoires
THE WEB MAGAZINE FOR GEOGRAPHIC INFORMATION SYSTEMS

PRACTICAL ASSESSMENT
NOVEMBER 2024

TerraLab AI Agent

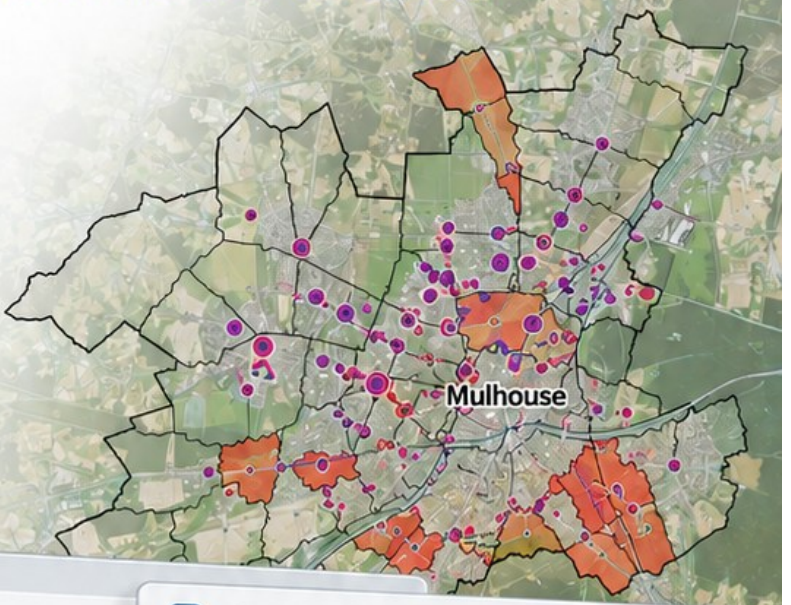
32 practical tests in QGIS 4.2

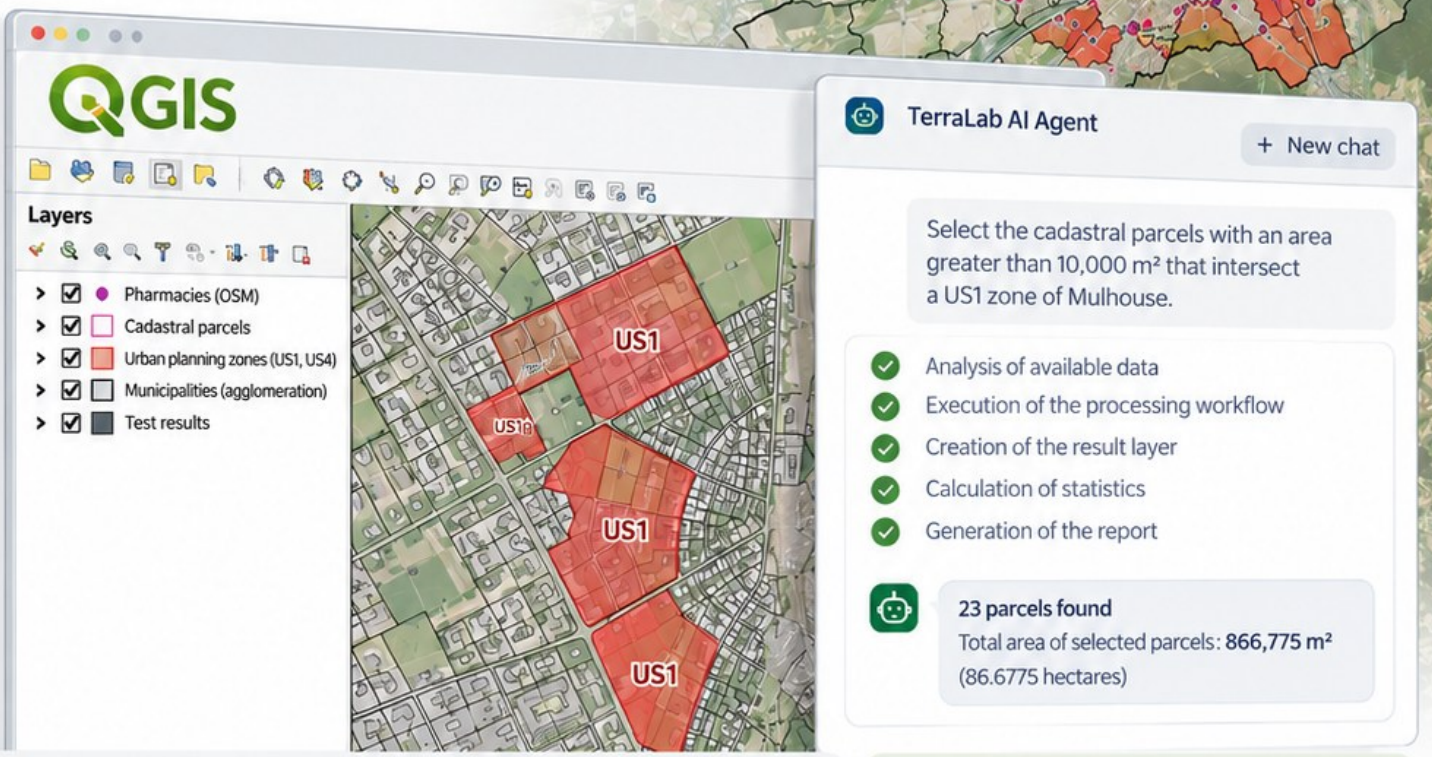
Prompts, executions, results and manual validation

32 detailed tests
Real-world use cases for spatial analysis, data management and cartography

Reproducible
Prompts, execution screenshots, QGIS results and agent reports

Verified results
Numerical values checked manually and additional controls





Contents of this document

1. Presentation of the 32 tests
2. For each test: prompt, execution screenshots,
2. results and agent report
3. Manual validation and observations
4. Resources to reproduce the tests

Data and project provided

- Initial QGIS project (.qgz)
- Two data layers
- Instructions for reproduction

Author: SIG & Territories
Evaluation conducted as part of a presentation of TerraLab AI Agent

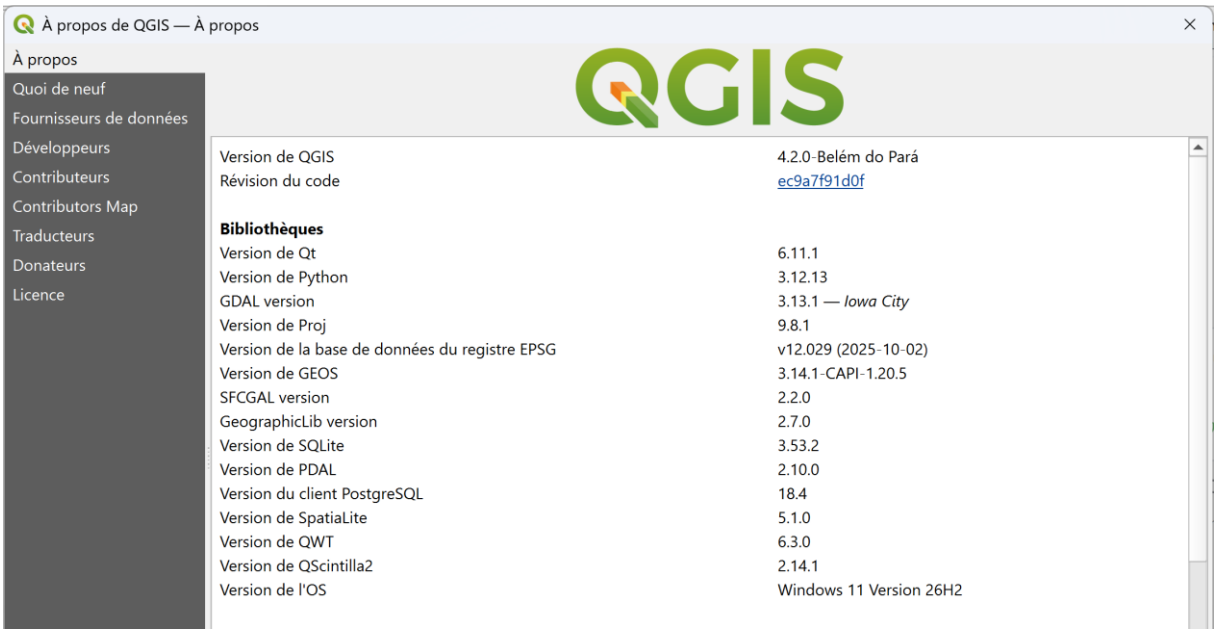
Software versions
QGIS 4.2
TerraLab AI Agent 1.9.2 (tested version)

Table of Contents

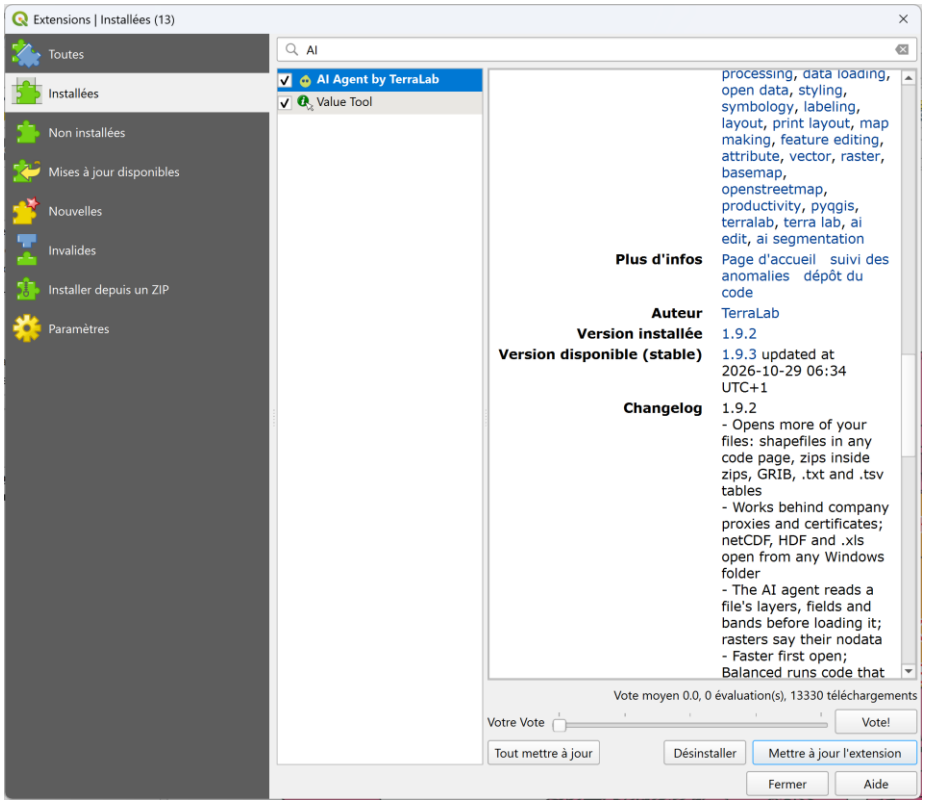
Tested Versions 3
Initial Project 4
Test 1: Simple, Reproducible, and Verifiable Test 4
Test 2: Run a Process 5
Test 3: Verify that the agent actually runs a QGIS process 7
Test 4: Request a Geometric Calculation 8
Test 5: Moving from Analysis to Mapping 11
Test 6: Comparing Total Population and Density 13
Test 7: Comparison of the Two Maps 16
Test 8: Ask AI Agent to Correct Editorial Errors 17
Test 9: Have the Agent Improve Their Own Map 20
Test 10: Deliberately Impossible Requestible 23
Test 11: Handling an Error During a Workflow 24
Test 12: Searching for OpenStreetMap Data 25
Test 13: Counting Pharmacies by Municipality 27
Test 14 — Retrieving Data from Another Source 30
Test 15 — Verify the Geometry and Consistency of the Result 32
Test 16 — Verify Zoning Through Spatial Analysis 34
Test 17 — Testing a Cadastral Search at the Municipal Level 35
Test 18 — Compare the Two Sources Without Loading the Cadastral Parcels 37
Test 19 — Load the Etalab Municipal Cadastral Database 39
Test 20 — Comprehensive Check of the Etalab Layer 43
Test 21 — Generate a Statistical Summary of the Cadastral Register 44
Test 22 — Extract Large Parcels and Verify the Results 47
Test 23 — Combining an Attribute-Based Criterion and a Spatial Criterion 50
Test 24 — Distinguishing Between Data Limitations and Geographic Results 52
Test 25 — Large Parcels Intersecting US1 53
Test 26 — Measure the proportion of each parcel located within US1 56
Test 27 — Distinguishing Significant Intersections from Marginal Overlaps 58
Test 28 — Reproducibility of a Spatial Analysis 60
Test 29 — Verify the Independence of the Calculation 62
Cross-validation with another GIS engine: Validation with PostGIS 64
Test 30 — Minimal prompt 65
Test 31 — Intermediate Prompt 67
Test 32 — Detailed Prompt 69

Tested Versions

QGis :4.2.0



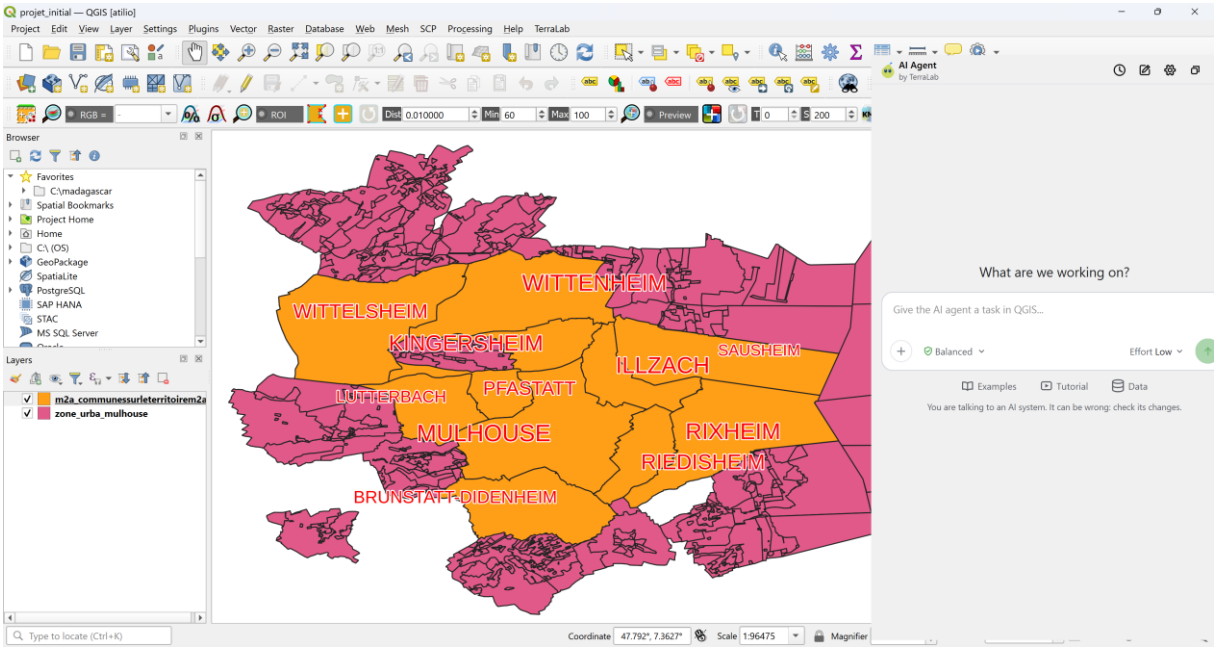
AI Agent plugin : 1.9.2



Initial Project

All tests are based on an initial project containing two layers for the Mulhouse metropolitan area:

- Cadastral parcels
- PLU zoning



Mapping Tests

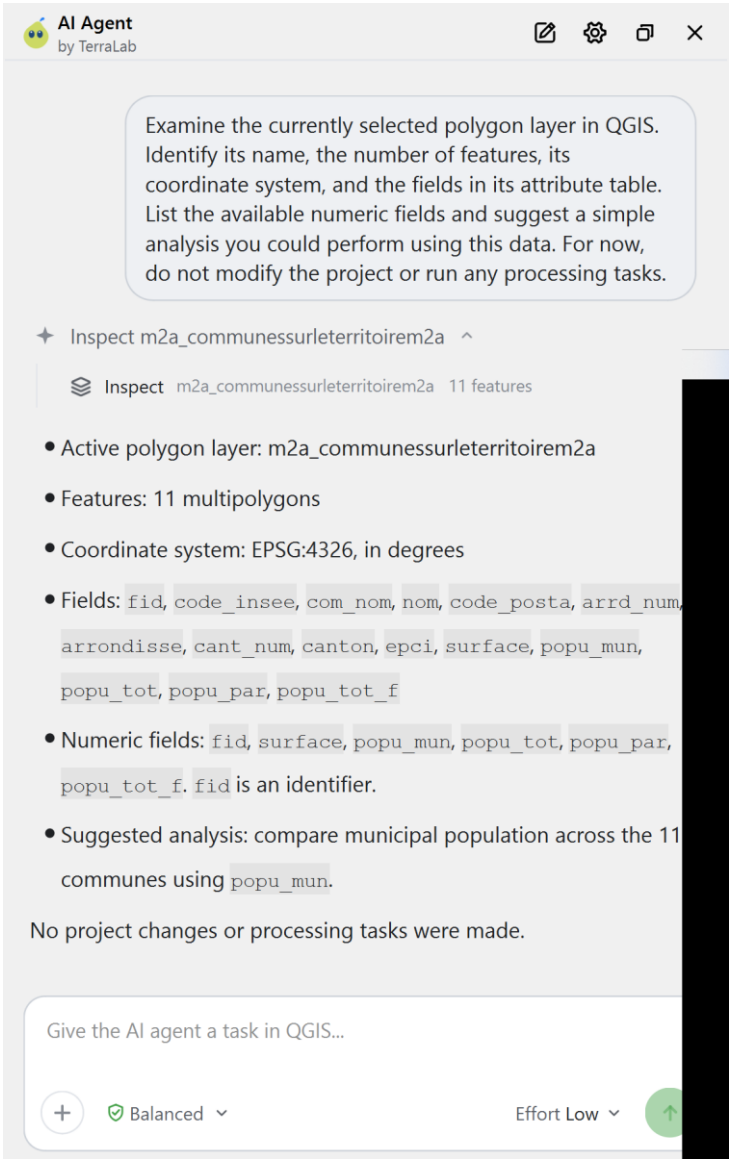
Test 1: Simple, Reproducible, and Verifiable Test

Objective: To perform a simple, reproducible, and verifiable test before asking AI Agent to carry out a complete workflow. The goal is to verify that it understands a geomatics instruction, actually performs the operations in QGIS 4.2, and produces a correct result.

Test prompt:

Examine the currently selected polygon layer in QGIS. Identify its name, the number of features, its coordinate system, and the fields in its attribute table. List the available numeric fields and suggest a simple analysis you could perform using this data. For now, do not modify the project or run any processing tasks.

Execution :



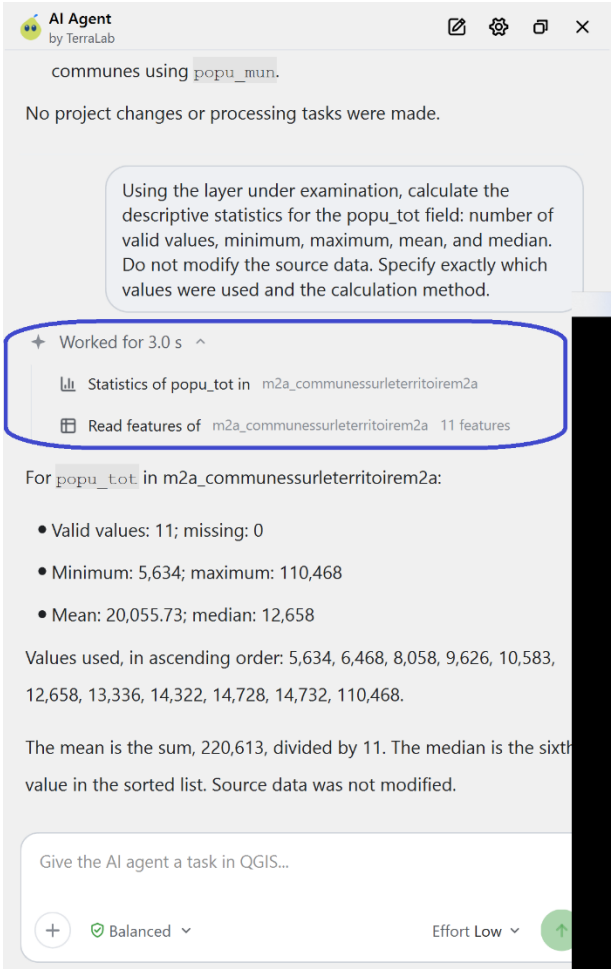
Test 2: Run a Process

Objective: Run a process

Test prompt:

Using the layer under examination, calculate the descriptive statistics for the popu_tot field: number of valid values, minimum, maximum, mean, and median. Do not modify the source data. Specify exactly which values were used and the calculation method.

Execution :



For popu_tot in m2a_communessurleterritoirem2a:

Valid values: 11; missing: 0

Minimum: 5,634; maximum: 110,468

Mean: 20,055.73; median: 12,658

Values used, in ascending order: 5,634, 6,468, 8,058, 9,626, 10,583, 12,658, 13,336, 14,322, 14,728, 14,732, 110,468.

The mean is the sum, 220,613, divided by 11. The median is the sixth value in the sorted list. Source data was not modified.

Validation :

What These First Two Tests Tell Us

They provide an initial assessment of three aspects of the agent's capabilities:

- Understanding of the project: Can the agent correctly access and interpret the information available in the project's layers?

- Data analysis: Can it identify the relevant attributes and perform the requested statistical calculations?
- Execution and reliability: Are the calculated values correct and verifiable?

These tests are not sufficient to evaluate the entire plugin, but they provide a useful baseline.

AI Agent correctly identified the layer context and responded to the statistical request. The reported values are arithmetically consistent and have been manually verified. The execution report also confirms that the source data was not modified.

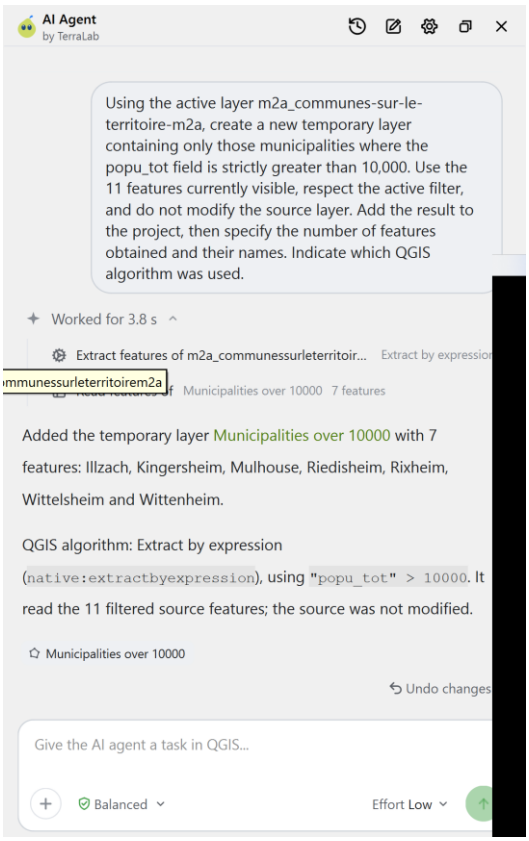
Test 3: Verify that the agent actually runs a QGIS process

Objective: Verify that the agent actually runs a QGIS process. We will ask it to create a new layer without modifying the original layer.

Test prompt:

Using the active layer m2a_communes-sur-le-territoire-m2a, create a new temporary layer containing only those municipalities where the popu_tot field is strictly greater than 10,000. Use the 11 features currently visible, respect the active filter, and do not modify the source layer. Add the result to the project, then specify the number of features obtained and their names. Indicate which QGIS algorithm was used.

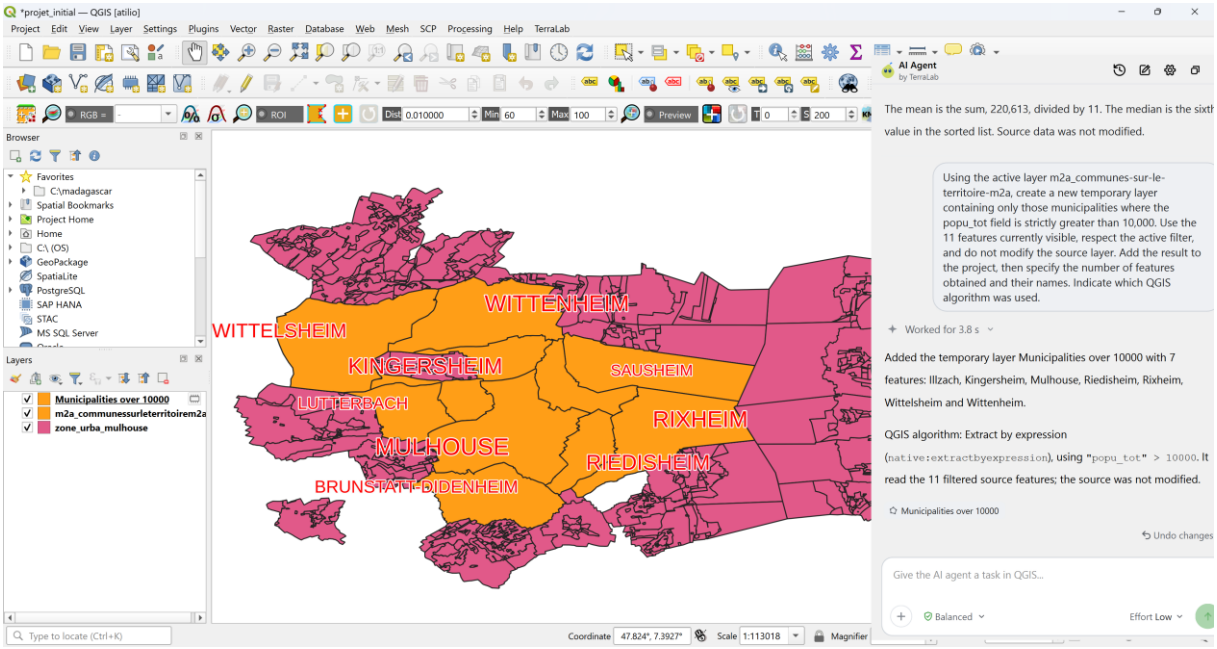
Execution :



Added the temporary layer Municipalities over 10000 with 7 features: Illzach, Kingersheim, Mulhouse, Riedisheim, Rixheim, Wittelsheim and Wittenheim.

QGIS algorithm: Extract by expression (`native:extractbyexpression`), using `"popu_tot" > 10000`. It read the 11 filtered source features; the source was not modified.

Result in QGIS:



Validation :

The execution confirms the use of `native:extractbyexpression`, the creation of a temporary layer containing seven features, and the preservation of the source layer.

This third test provides additional evidence of the agent's capabilities: AI Agent did not merely generate a filtering expression; it actually executed an algorithm from the QGIS Processing framework and added the resulting layer to the project.

The seven extracted municipalities correspond to the requested population criterion, and the source data remained unchanged.

.

Test 4: Request a Geometric Calculation

Objective: Request a geometric calculation: calculate the area of the municipalities in hectares, taking the coordinate system into account.

The layer is in EPSG:4326, so the coordinates are expressed in degrees. A simple calculation of area based on the coordinates would be incorrect. The agent must choose an appropriate method.

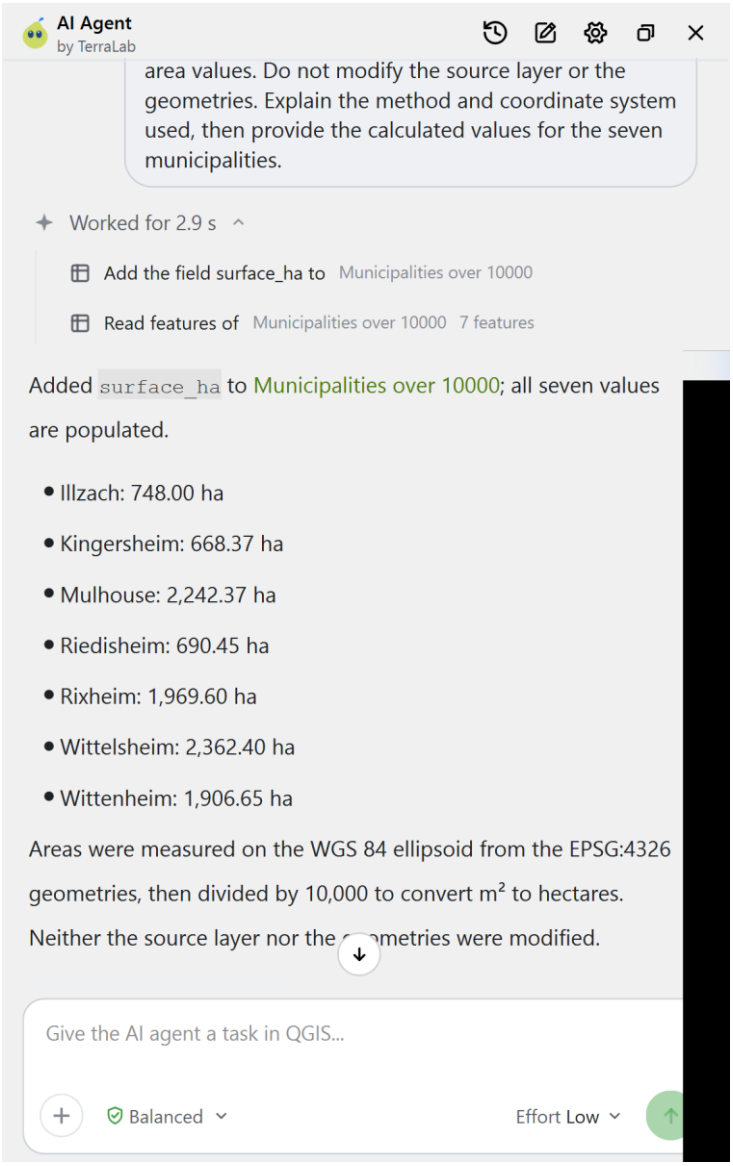
What we will observe

- Does the user understand that degrees are not units of area?
- Does the user choose a geodetic calculation or a projection suitable for the Mulhouse region?
- Is the field actually created and populated in the attribute table?
- Are the results plausible and reproducible?

Test prompt:

On the temporary layer "Municipalities with more than 10,000 inhabitants," add a numeric field named `surface_ha` containing the area of each municipality in hectares. The layer is currently in EPSG:4326. Choose an appropriate geometric method to obtain reliable area values. Do not modify the source layer or the geometries. Explain the method and coordinate system used, then provide the calculated values for the seven municipalities.

Execution :

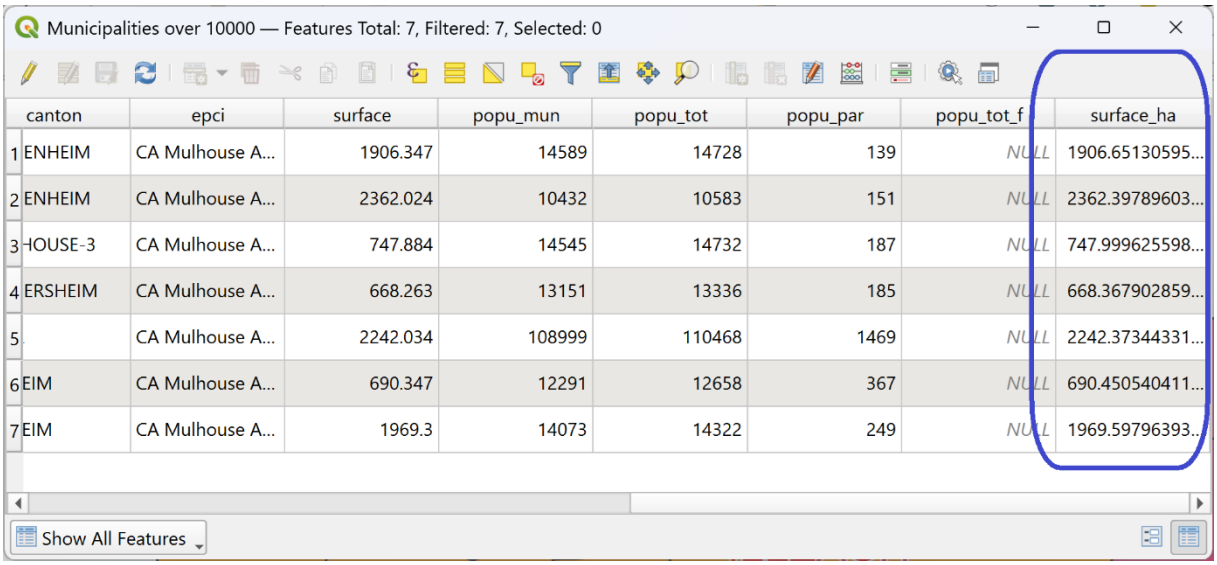


Added `surface_ha` to Municipalities over 10000; all seven values are populated.

- Illzach: 748.00 ha
- Kingersheim: 668.37 ha
- Mulhouse: 2,242.37 ha
- Riedisheim: 690.45 ha
- Rixheim: 1,969.60 ha
- Wittelsheim: 2,362.40 ha
- Wittenheim: 1,906.65 ha

Areas were measured on the WGS 84 ellipsoid from the EPSG:4326 geometries, then divided by 10,000 to convert m² to hectares. Neither the source layer nor the geometries were modified.

Result in QGis:



	canton	epci	surface	popu_mun	popu_tot	popu_par	popu_tot_f	surface_ha
1	ENHEIM	CA Mulhouse A...	1906.347	14589	14728	139	NULL	1906.65130595...
2	ENHEIM	CA Mulhouse A...	2362.024	10432	10583	151	NULL	2362.39789603...
3	HOUSE-3	CA Mulhouse A...	747.884	14545	14732	187	NULL	747.999625598...
4	ERSHEIM	CA Mulhouse A...	668.263	13151	13336	185	NULL	668.367902859...
5		CA Mulhouse A...	2242.034	108999	110468	1469	NULL	2242.37344331...
6	EIM	CA Mulhouse A...	690.347	12291	12658	367	NULL	690.450540411...
7	EIM	CA Mulhouse A...	1969.3	14073	14322	249	NULL	1969.59796393...

Validation:

AI Agent added the `surface_ha` field to the temporary layer and populated it with the calculated areas, measured on the WGS 84 ellipsoid. The original geometries and source layer were preserved.

This is an appropriate approach for calculating areas from geometries stored in EPSG:4326, avoiding direct planimetric area calculations in geographic coordinates. A geodetic area calculation can be performed while retaining the layer's original CRS.

The seven values were correctly populated in the attribute table and manually verified.

According to the QGIS documentation, the `$area` expression performs an ellipsoidal area calculation based on the ellipsoid and measurement units configured in the project. The execution report confirms the use of the WGS 84 ellipsoid, although it does not specify which expression or function was used.

Test 5: Moving from Analysis to Mapping

Objective: Move from analysis to mapping: verify whether AI Agent can also produce a usable map, rather than just performing calculations.

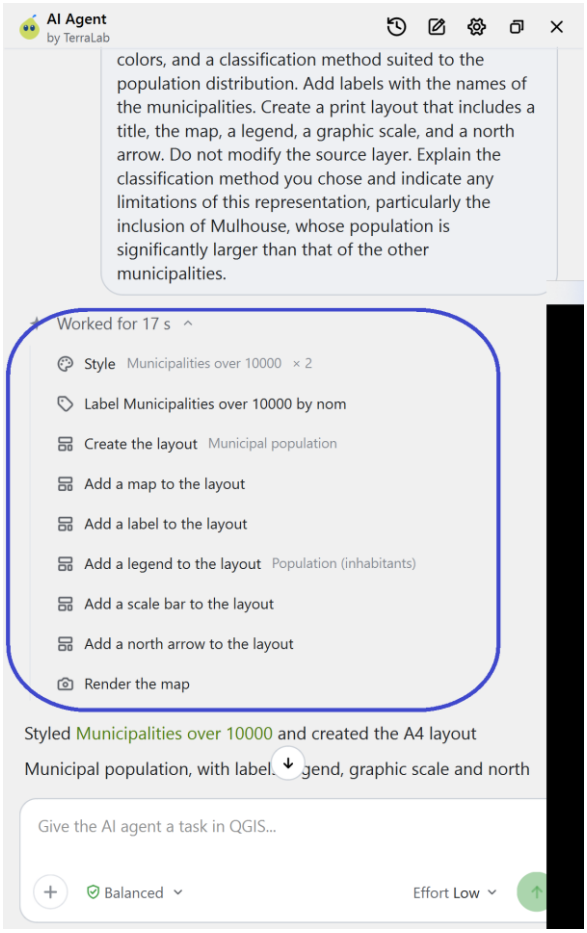
What this test will tell us

- Can the agent choose an appropriate classification for highly skewed data?
- Does it actually configure the symbology and labels in QGIS?
- Does it create a complete layout with functional elements?
- Is the result readable, or merely technically correct?

Test prompt:

Using the temporary layer "Municipalities with more than 10,000 inhabitants," create a map representation of the total population (popu_tot). Use graduated symbology with five classes, progressively darker colors, and a classification method suited to the population distribution. Add labels with the names of the municipalities. Create a print layout that includes a title, the map, a legend, a graphic scale, and a north arrow. Do not modify the source layer. Explain the classification method you chose and indicate any limitations of this representation, particularly the inclusion of Mulhouse, whose population is significantly larger than that of the other municipalities.

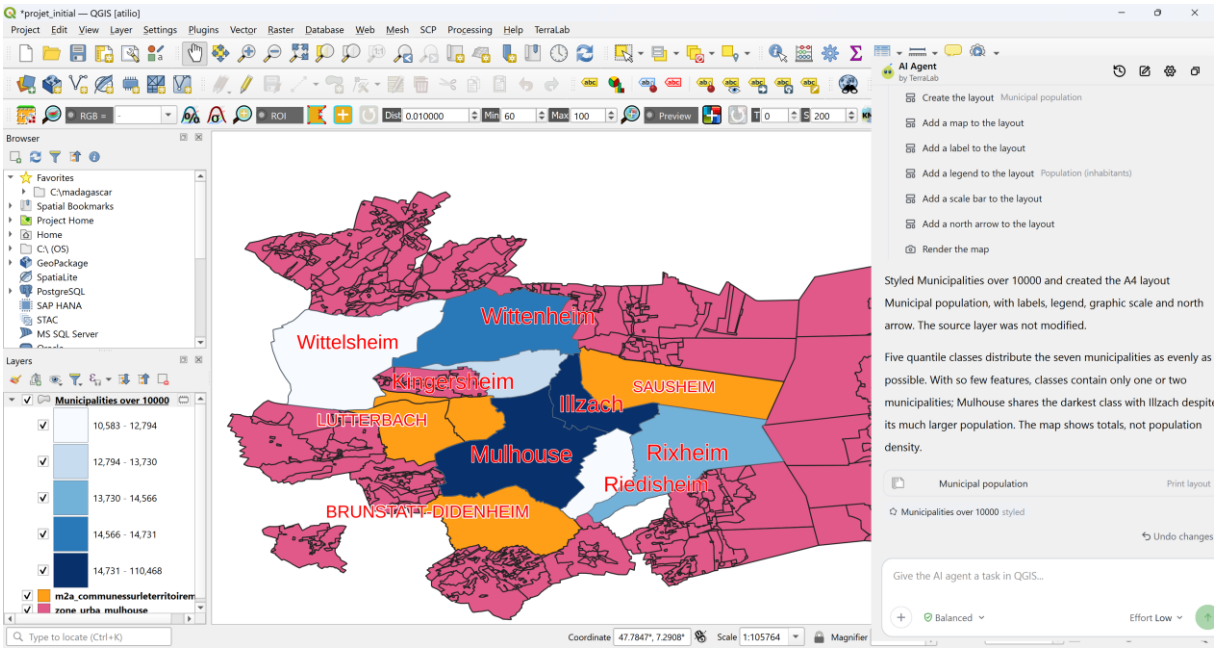
Execution :

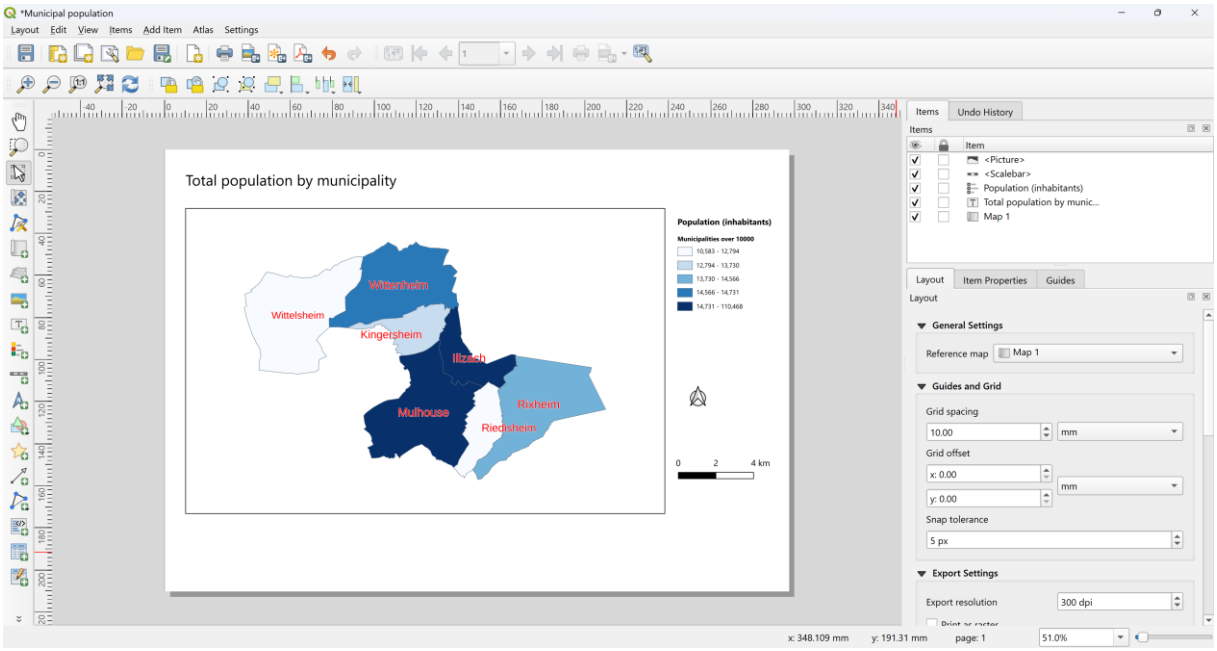


Styled Municipalities over 10000 and created the A4 layout Municipal population, with labels, legend, graphic scale and north arrow. The source layer was not modified.

Five quantile classes distribute the seven municipalities as evenly as possible. With so few features, classes contain only one or two municipalities; Mulhouse shares the darkest class with Illzach despite its much larger population. The map shows totals, not population density..

Result in QGIS:





Validation:

AI Agent did not simply apply a symbology; it created an A4 map layout and provided an observation about the limitations of the chosen classification method.

The map uses five quantile classes to distribute the seven municipalities as evenly as possible. With so few features, each class contains only one or two municipalities. Mulhouse, with 110,468 inhabitants, shares the darkest class with Illzach, despite its much larger population.

This illustrates a limitation of quantile classification when applied to a dataset containing very few features: municipalities with substantially different values may end up in the same class.

The agent also correctly points out that the map represents total population, not population density.

The layout itself still needs to be examined to assess the readability of labels, the legend, the graphic scale and the map extent. The execution report alone is not sufficient to assess the overall cartographic quality.

Test 6: Comparing Total Population and Density

Objective: Compare total population and density. The agent refers to a density map. This is a good follow-up, provided the density is calculated using a consistent area. Density is calculated by dividing the population by the land area. For Mulhouse, using the values provided earlier, the result is approximately 4,927 inhabitants per square kilometer. This calculation is for illustrative purposes only: the exact area and population figures from the data must be verified.

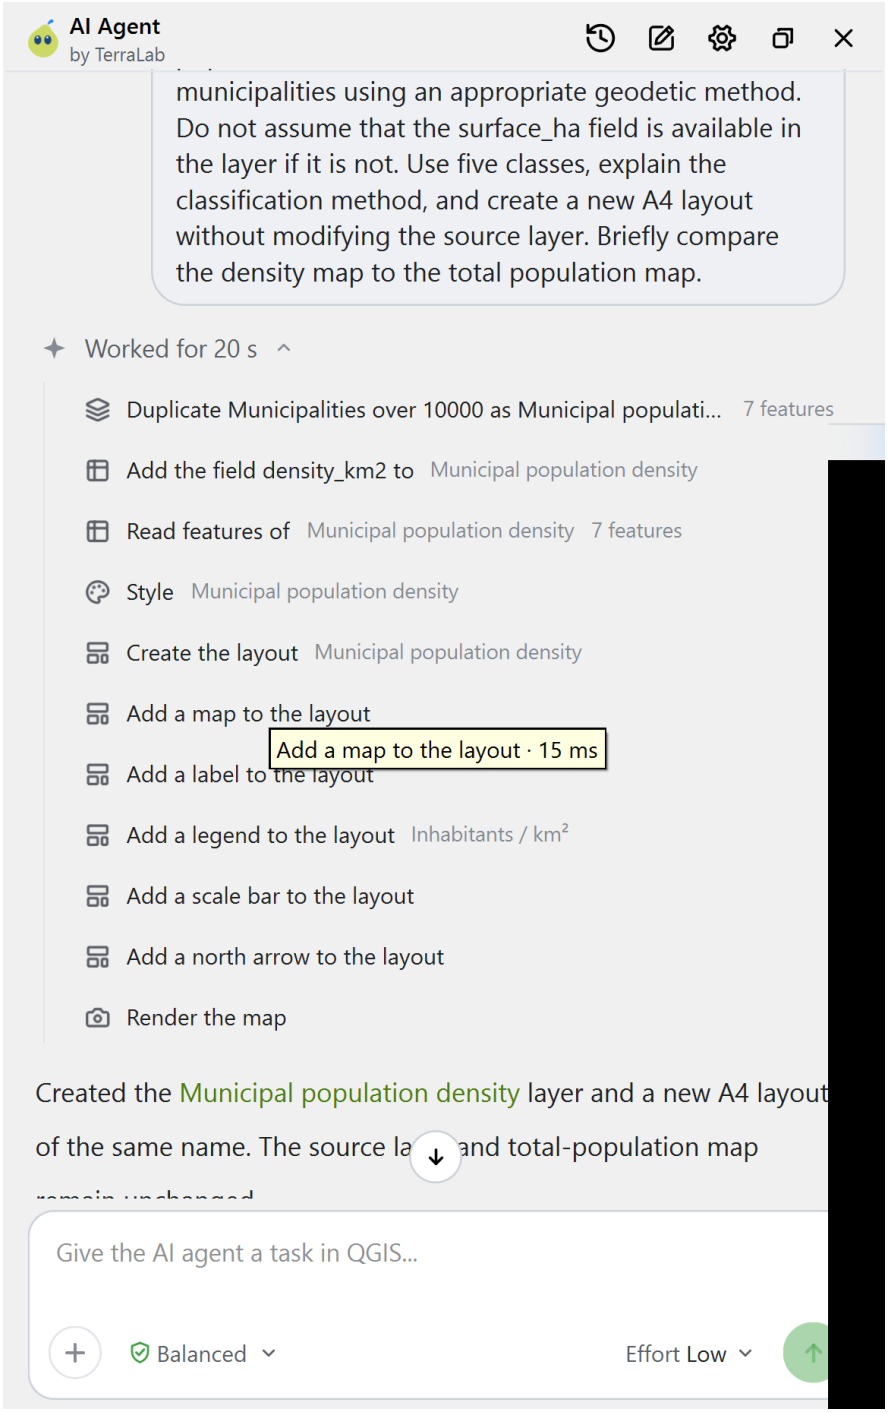
What this test will assess

- The correct application of results from the previous test.
- The distinction between total population and density.
- The consistency of units and geodetic calculations.
- The ability to create a second layout without overwriting the first.

Test prompt:

Using the "Municipalities with more than 10,000 inhabitants" layer, create a choropleth map showing population density in inhabitants per km². Use the popu_tot field and calculate the area of the municipalities using an appropriate geodetic method. Do not assume that the surface_ha field is available in the layer if it is not. Use five classes, explain the classification method, and create a new A4 layout without modifying the source layer. Briefly compare the density map to the total population map.

Execution :

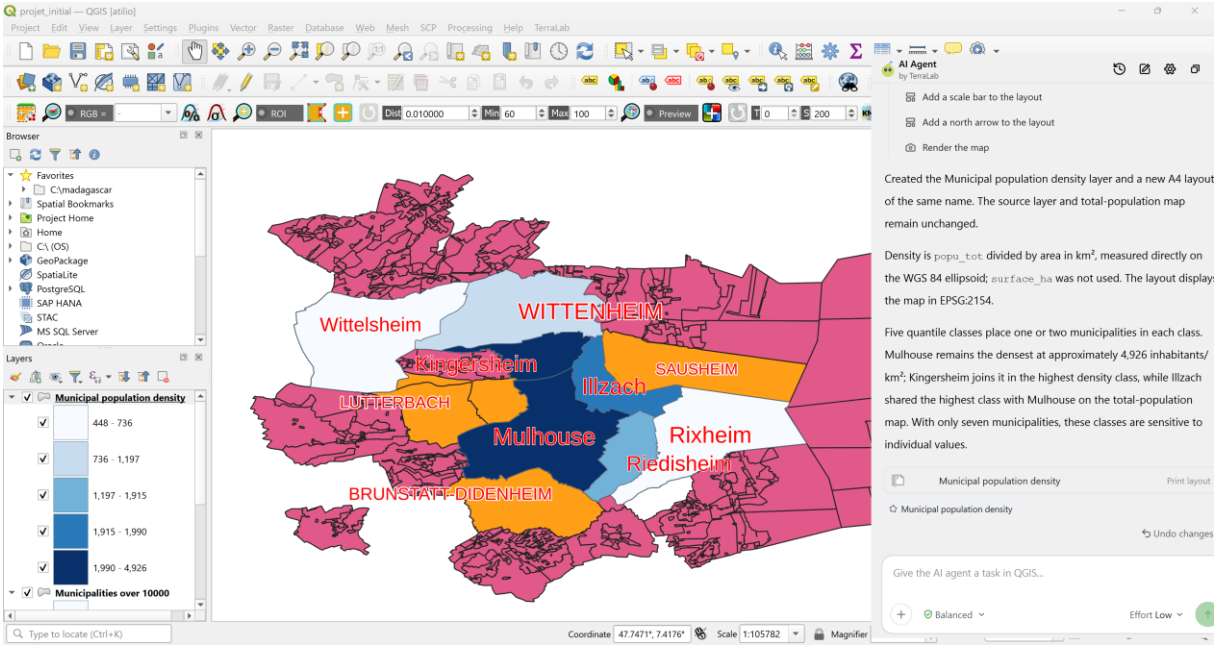


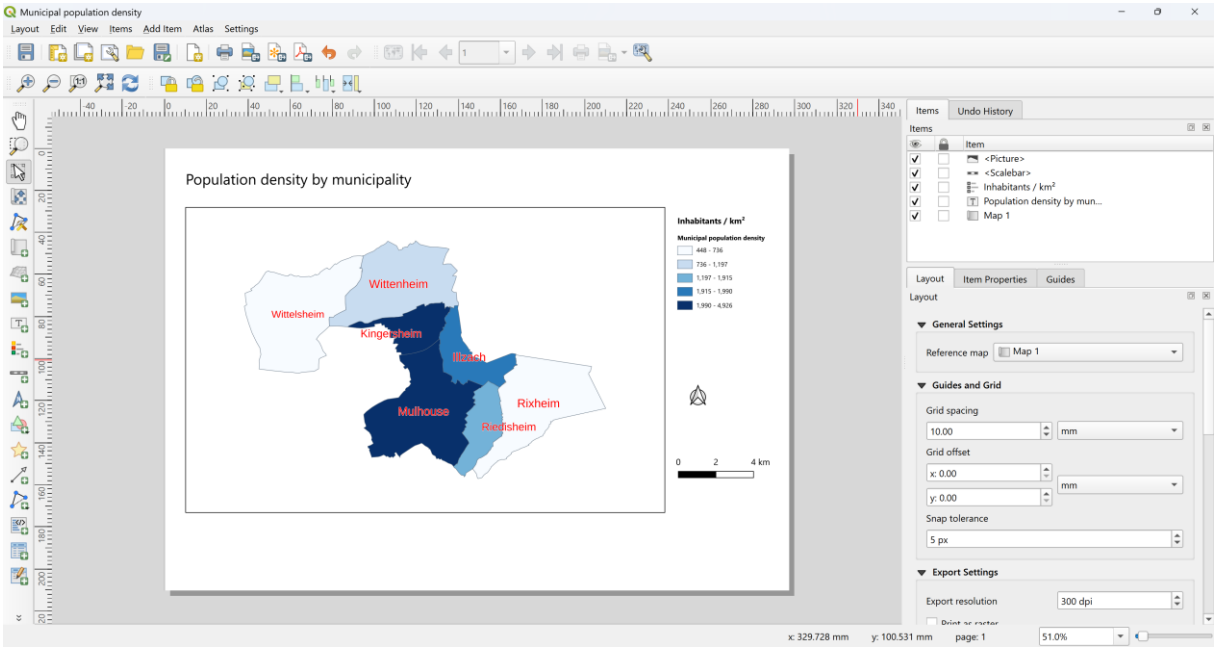
Created the Municipal population density layer and a new A4 layout of the same name. The source layer and total-population map remain unchanged.

Density is `popu_tot` divided by area in km², measured directly on the WGS 84 ellipsoid; `surface_ha` was not used. The layout displays the map in EPSG:2154.

Five quantile classes place one or two municipalities in each class. Mulhouse remains the densest at approximately 4,926 inhabitants/km²; Kingersheim joins it in the highest density class, while Illzach shared the highest class with Mulhouse on the total-population map. With only seven municipalities, these classes are sensitive to individual values

Result in QGIS





Validation:

AI Agent performed the geodetic area calculation using the WGS 84 ellipsoid, created a derived population density layer, and generated a second A4 layout. The original source layer and the total-population map remained unchanged.

The density calculation was based on `popu_tot` divided by the municipality area in square kilometres, measured directly on the ellipsoid. The previously calculated `surface_ha` field was not reused.

The resulting map illustrates the difference between total population and population density. Mulhouse remains the densest municipality, at approximately 4,926 inhabitants/km², while Kingersheim joins it in the highest density class. This differs from the total-population map, where Illzach shared the highest class with Mulhouse.

The five quantile classes contain only one or two municipalities each. As the agent correctly notes, classifications based on such a small number of features are sensitive to individual values.

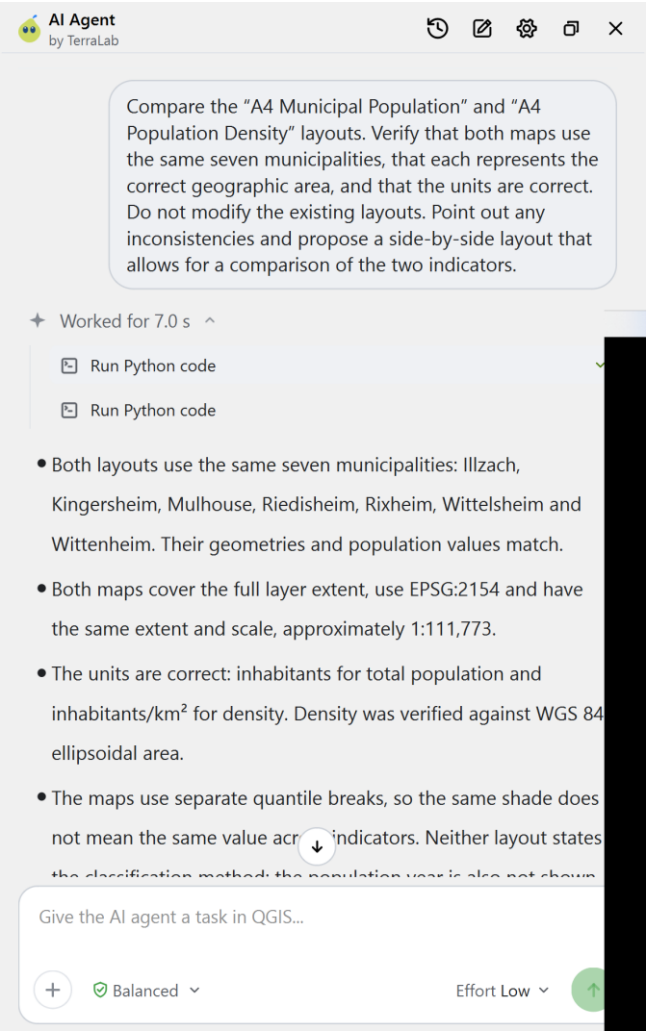
Test 7: Comparison of the Two Maps

Objective: To compare the two maps. This test serves two purposes: it verifies whether the agent can locate and compare objects already created in the project, and whether the agent can assess the consistency of a cartographic analysis.

Test prompt:

Compare the "A4 Municipal Population" and "A4 Population Density" layouts. Verify that both maps use the same seven municipalities, that each represents the correct geographic area, and that the units are correct. Do not modify the existing layouts. Point out any inconsistencies and propose a side-by-side layout that allows for a comparison of the two indicators.

Execution :



Both layouts use the same seven municipalities: Illzach, Kingersheim, Mulhouse, Riedisheim, Rixheim, Wittelsheim and Wittenheim. Their geometries and population values match.

Both maps cover the full layer extent, use EPSG:2154 and have the same extent and scale, approximately 1:111,773.

The units are correct: inhabitants for total population and inhabitants/km² for density. Density was verified against WGS 84 ellipsoidal area.

The maps use separate quantile breaks, so the same shade does not mean the same value across indicators. Neither layout states the classification method; the population year is also not shown.

For a side-by-side layout, use two equal map frames with the same extent and scale, separate titled legends with units, and a shared scale bar. No existing layout was modified.

Validation :

AI Agent did more than compare the two layers: it checked their consistency, compared the map layouts and identified several points requiring attention. This demonstrates its usefulness for reviewing cartographic outputs.

Several observations are particularly relevant:

Data currency: The population year is not displayed. Without this information, readers may assume that the maps represent the current situation, even if the underlying data are older.

Classification method: Neither layout identifies the classification method used.

Comparison of colours: The two maps represent different indicators and use separate quantile breaks. Consequently, the same colour does not represent an equivalent value on both maps.

Cartographic consistency: Both maps use the same seven municipalities, the same extent and scale, and the appropriate measurement units.

The agent also proposed a practical solution for a side-by-side presentation: two equal map frames with the same extent and scale, separate legends with titles and units, and a shared scale bar.

Test 8: Ask AI Agent to Correct Editorial Errors

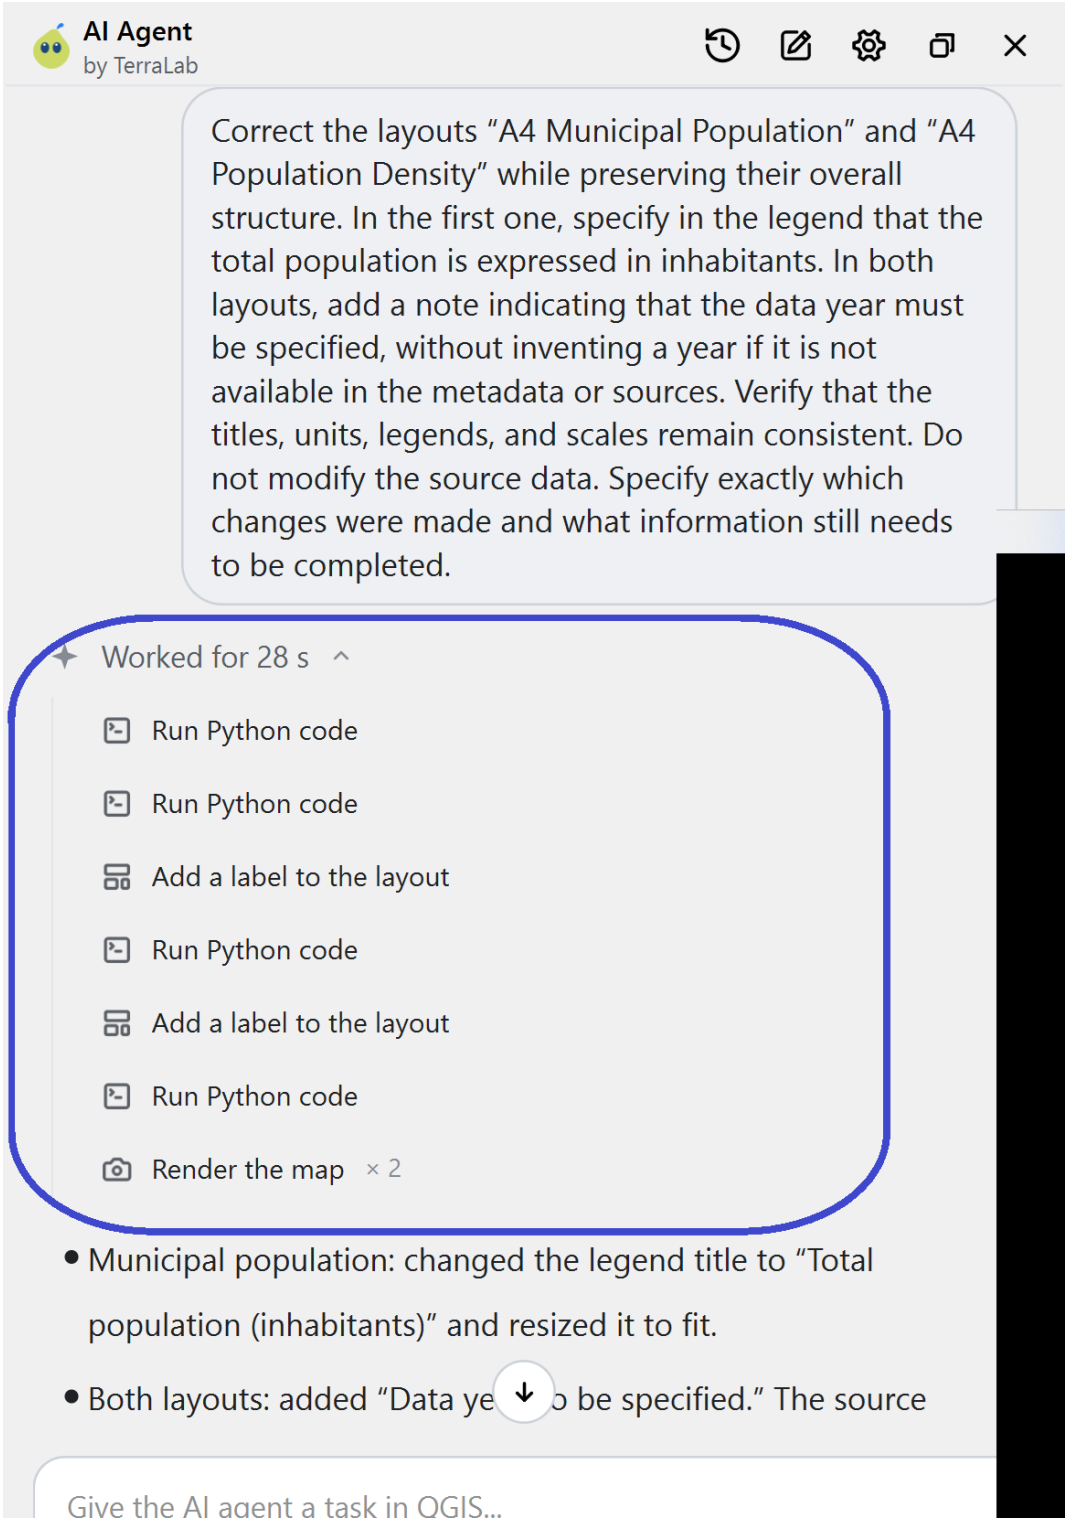
Objective: Ask AI Agent to correct editorial errors. Before moving on to a more complex workflow, we will ask it to correct the two existing layouts. This will test its ability to modify existing objects without having to rebuild them entirely.

Test prompt:

Correct the layouts “A4 Municipal Population” and “A4 Population Density” while preserving their overall structure. In the first one, specify in the legend that the total population is expressed in inhabitants. In both layouts, add a note indicating that the data year must be specified, without

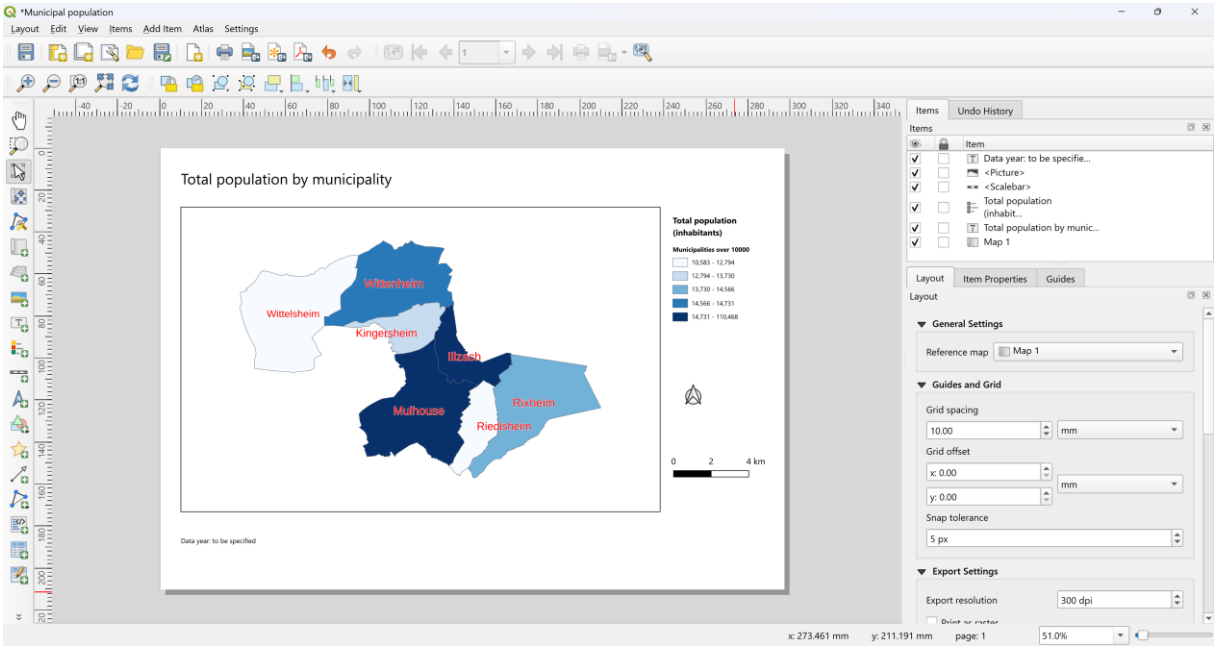
inventing a year if it is not available in the metadata or sources. Verify that the titles, units, legends, and scales remain consistent. Do not modify the source data. Specify exactly which changes were made and what information still needs to be completed.

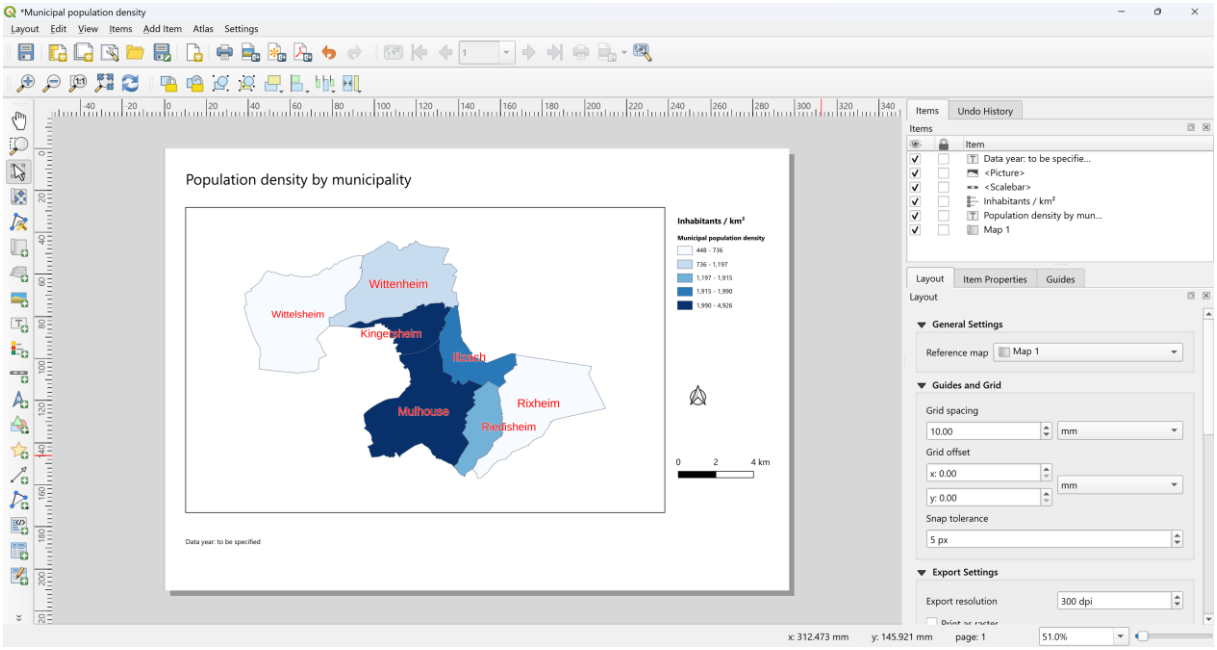
Execution :



- Municipal population: changed the legend title to “Total population (inhabitants)” and resized it to fit.
- Both layouts: added “Data year: to be specified.” The source layer’s metadata did not identify a year; it still needs to be confirmed.
- Checked both titles, legends and units. Both retain the same map scale and 0–4 km graphic scale. Source data was not modified

Result in QGIS:





Validation :

AI Agent demonstrated its ability to edit existing layouts and apply specific cartographic improvements in response to feedback.

It changed the population legend title to "Total population (inhabitants)" and resized it to fit. It also added a placeholder for the data year to both layouts, correctly acknowledging that the source metadata did not provide this information.

However, these corrections do not necessarily guarantee optimal cartographic quality. A visual review of the layouts remains useful to assess the overall balance of the page, the readability of municipality labels and the formatting of legend values.

The agent also preserved the existing map scale and graphic scale, and did not modify the source data.

Test 9: Have the Agent Improve Their Own Map

Objective: An interesting additional test: have the agent improve their own map

Test prompt:

Improve the “A4 Population Density” layout while keeping the map, the data, and the five existing categories.

Enlarge the map to make better use of the available space on the A4 page.

Slightly reduce the size of the labels and choose a color that remains legible without overwhelming the map.

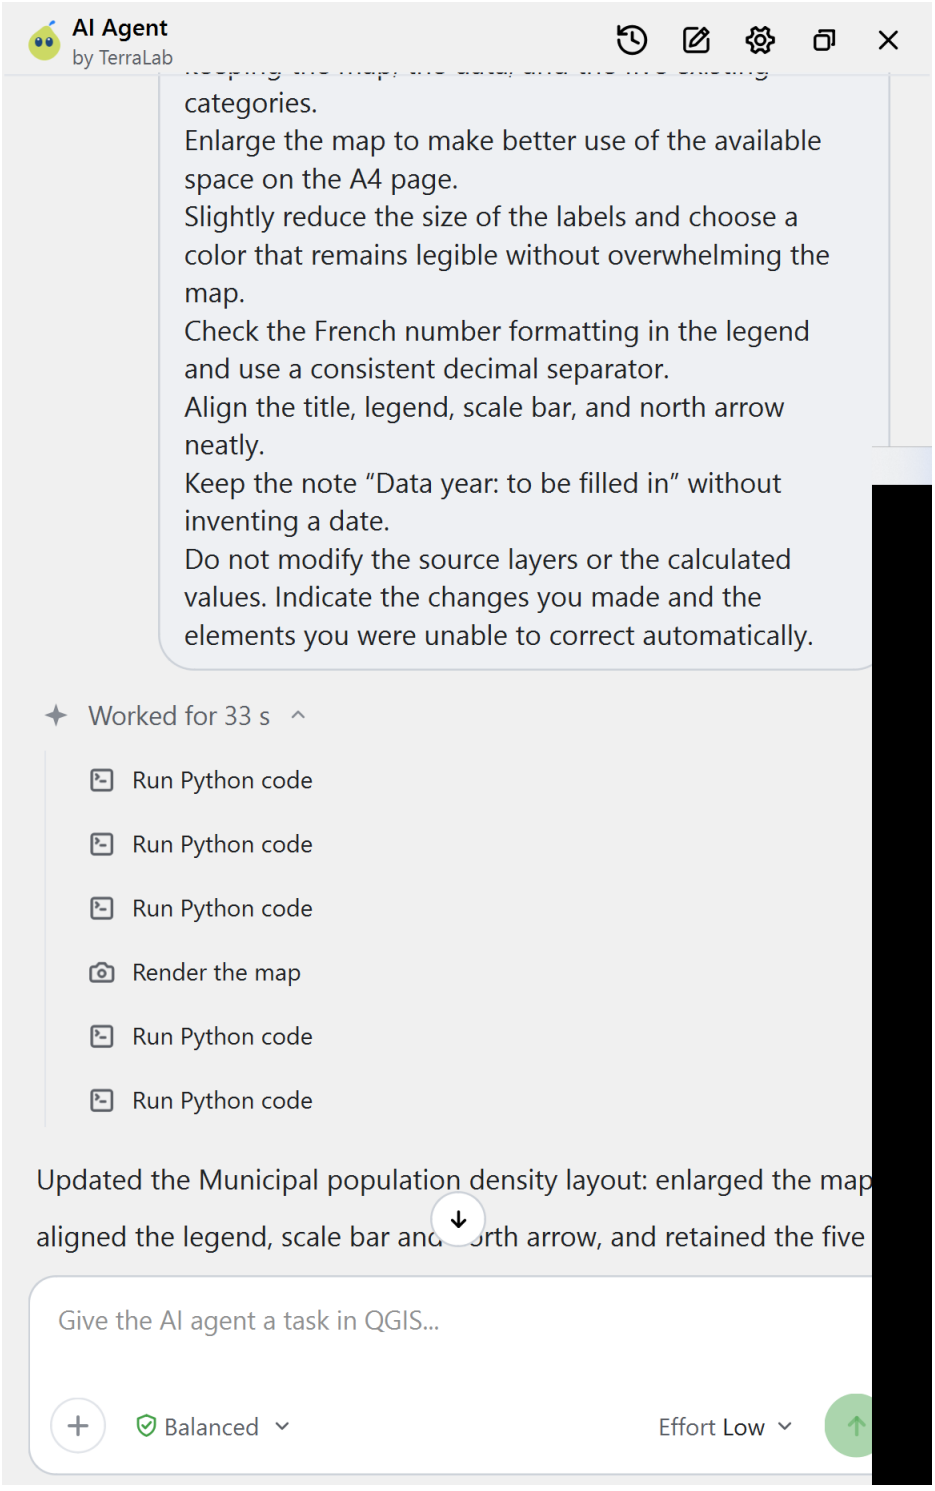
Check the French number formatting in the legend and use a consistent decimal separator.

Align the title, legend, scale bar, and north arrow neatly.

Keep the note “Data year: to be filled in” without inventing a date.

Do not modify the source layers or the calculated values. Indicate the changes you made and the elements you were unable to correct automatically.

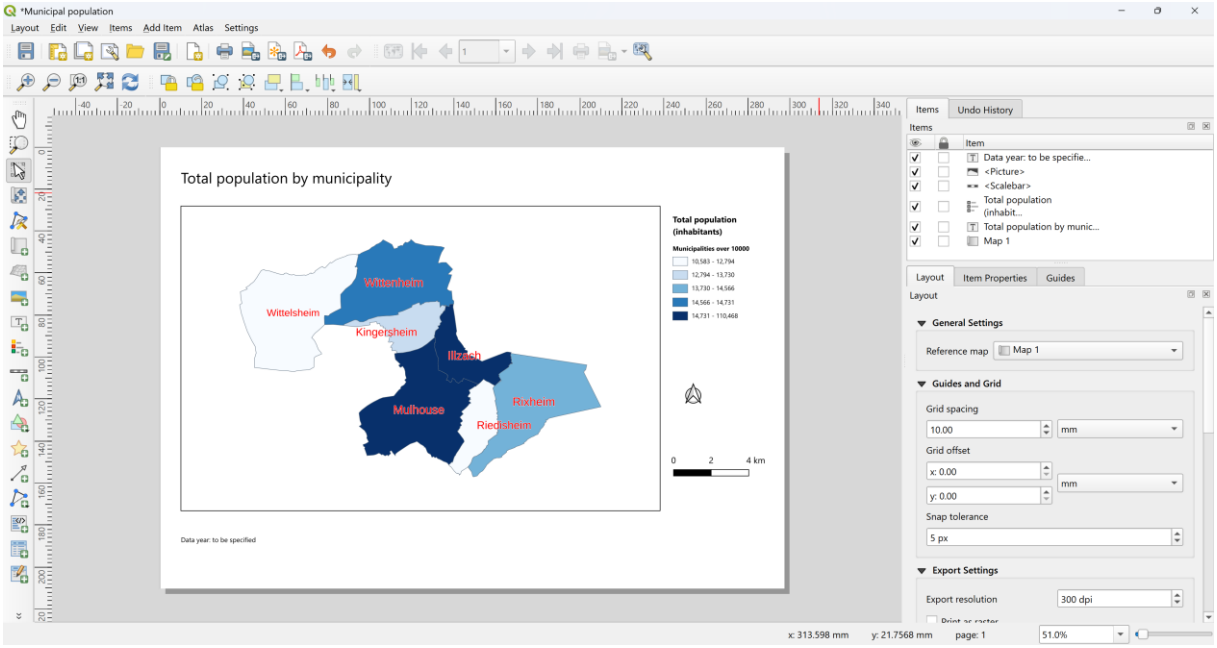
Execution :

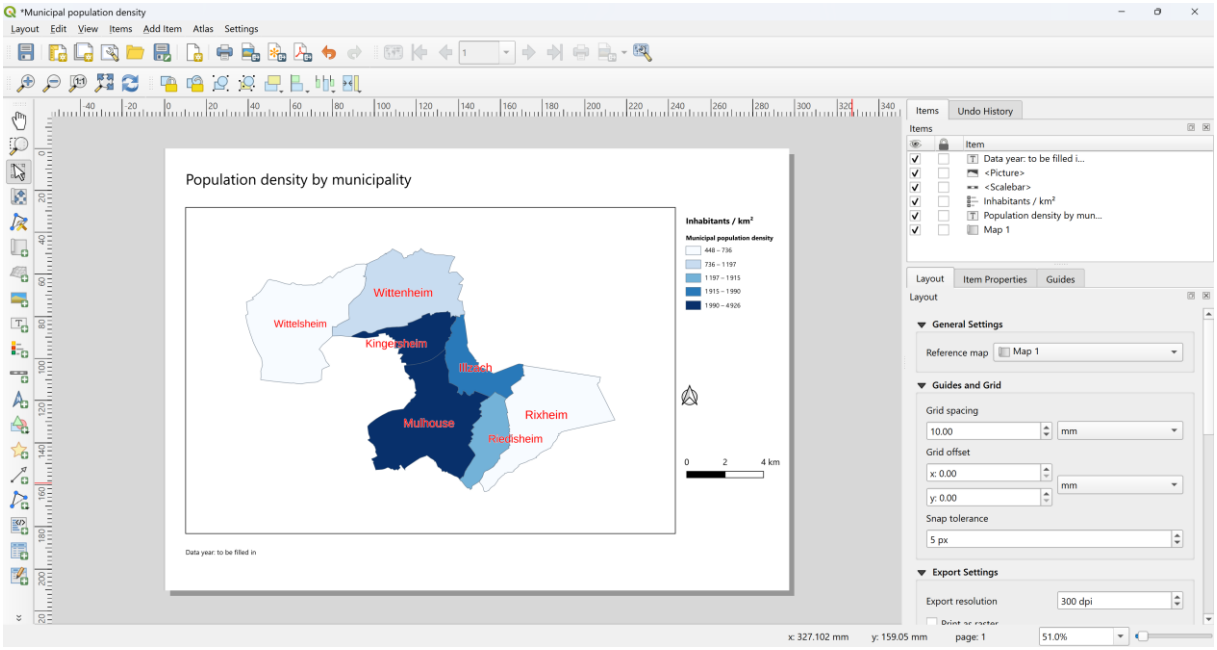


Updated the Municipal population density layout: enlarged the map, aligned the legend, scale bar and north arrow, and retained the five classes. Legend values use spaced thousands and no decimals. The note now reads “Data year: to be filled in.”

A layout-only 8 pt dark label style was set, but the preview still showed prominent red labels, so the label correction could not be verified. The enlarged map’s scale is now approximately 1:109,374. Source layers and calculated values were not changed.

Result in QGIS:





Updated the Municipal population density layout: enlarged the map, aligned the legend, scale bar and north arrow, and retained the five classes. Legend values use spaced thousands and no decimals. The note now reads “Data year: to be filled in.”

A layout-only 8 pt dark label style was set, but the preview still showed prominent red labels, so the label correction could not be verified. The enlarged map’s scale is now approximately 1:109,374. Source layers and calculated values were not changed.

Validation:

AI Agent successfully followed the layout instructions. The map is significantly larger, and the legend, scale bar and north arrow have been rearranged. The five classification classes were retained, and the legend values are now displayed without decimal places.

The agent also attempted to apply a smaller, dark label style. However, the preview still showed prominent red labels, so this particular correction could not be verified.

This test provides another example of iterative map improvement through dialogue with the agent. It also confirms that layout changes can be made while preserving the source layers and calculated values.

What We've Demonstrated So Far

The first nine tests have covered a range of operations, from identifying and analysing layers to creating, comparing and refining map layouts. One of the most interesting capabilities demonstrated so far is the agent's ability to modify an existing layout in response to specific feedback, while preserving the underlying data and classification.

Test 10: Deliberately Impossible Requestible

Objective: A deliberately impossible request. This test checks whether the agent recognizes an error, avoids fabricating an average, and clearly explains what is missing.

Test prompt:

In the "Municipalities with more than 10,000 residents" layer, calculate the average of the field `champ_inexistant`, which does not exist in the attribute table. Do not create any fields or modify any data. Explain the problem encountered and propose a solution, without inventing a result.

Execution :

`champ_inexistant` is not a field in Municipalities over 10000, so its average cannot be calculated.

Use an existing numeric field, such as `popu_tot`, or obtain data containing the intended field. No fields or data were modified.

Validation :

AI Agent explicitly identifies the missing field, does not invent a result, and suggests alternative solutions.

The response is appropriate. However, there is one limitation: since the request could be rejected as soon as the field name was checked, this test does not yet show how the agent handles an error occurring midway through a multi-step workflow involving several operations.

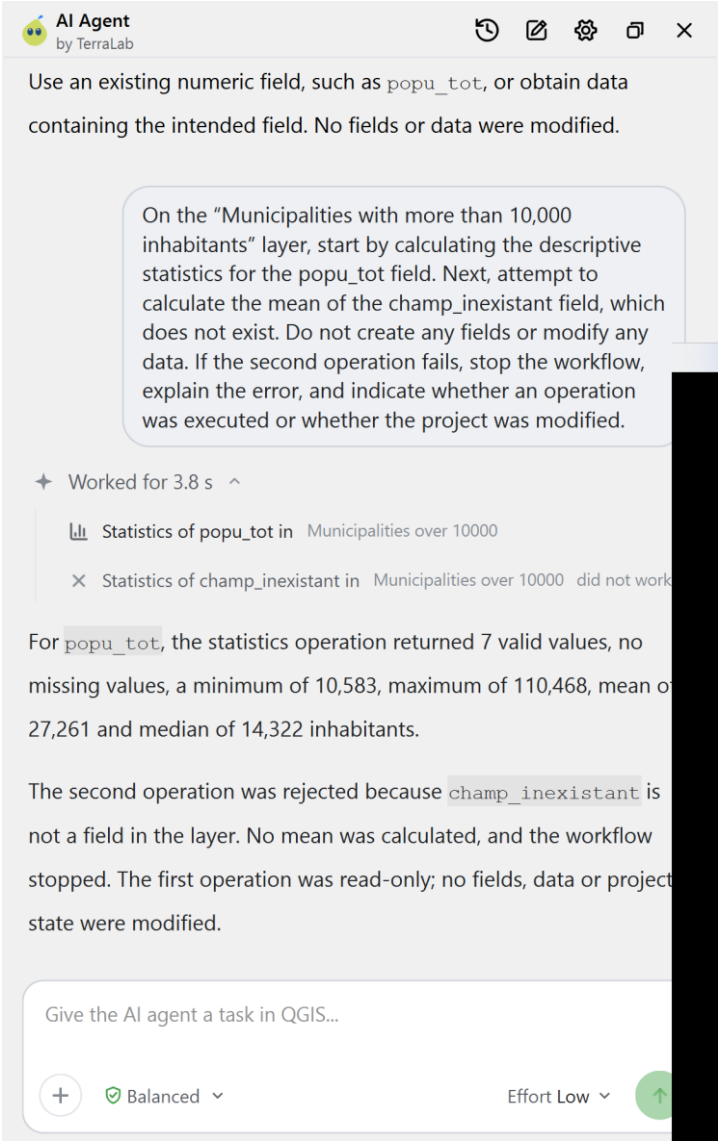
Test 11: Handling an Error During a Workflow

Objective: Handling an error during a workflow: trigger a sequence of two operations, the second of which is impossible. The goal is to observe whether the workflow terminates properly or leaves a project in a partially modified state.

Test prompt:

On the "Municipalities with more than 10,000 inhabitants" layer, start by calculating the descriptive statistics for the popu_tot field. Next, attempt to calculate the mean of the champ_inexistant field, which does not exist. Do not create any fields or modify any data. If the second operation fails, stop the workflow, explain the error, and indicate whether an operation was executed or whether the project was modified.

Execution :



For `popu_tot`, the statistics operation returned 7 valid values, no missing values, a minimum of 10,583, maximum of 110,468, mean of 27,261 and median of 14,322 inhabitants.

The second operation was rejected because `champ_inexistant` is not a field in the layer. No mean was calculated, and the workflow stopped. The first operation was read-only; no fields, data or project state were modified.

Validation :

The plugin successfully executed the first operation, calculating statistics for the seven valid values in `popu_tot`. It then encountered the expected error during the second operation because `champ_inexistant` does not exist in the layer.

The workflow stopped at that point, without attempting to calculate a mean for the missing field and without modifying any data or the project state.

This test demonstrates that the agent can execute a multi-step workflow, detect an error in a subsequent operation, and stop without reporting a result that was not calculated.

Tests Searching for Geographic Data

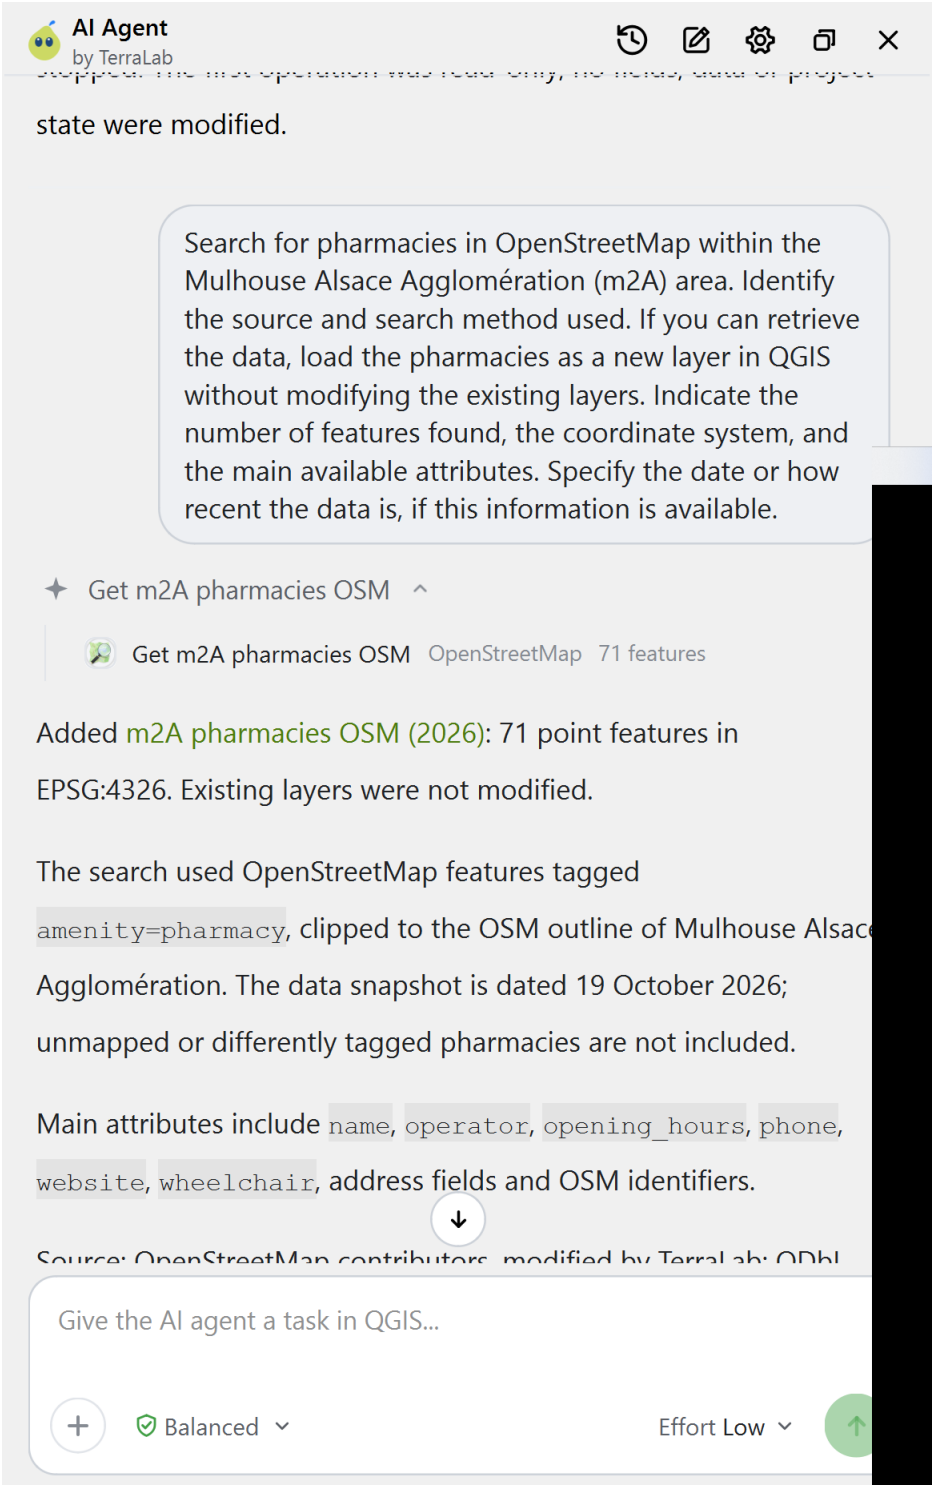
Test 12: Searching for OpenStreetMap Data

Objective: Search for OpenStreetMap data. To stay within the context of the project on the municipalities of Mulhouse Alsace Agglomération, let's start with pharmacies. This is a relatively simple geographic feature that can be identified in OpenStreetMap**.**

Test prompt:

Search for pharmacies in OpenStreetMap within the Mulhouse Alsace Agglomération (m2A) area. Identify the source and search method used. If you can retrieve the data, load the pharmacies as a new layer in QGIS without modifying the existing layers. Indicate the number of features found, the coordinate system, and the main available attributes. Specify the date or how recent the data is, if this information is available.

Execution :



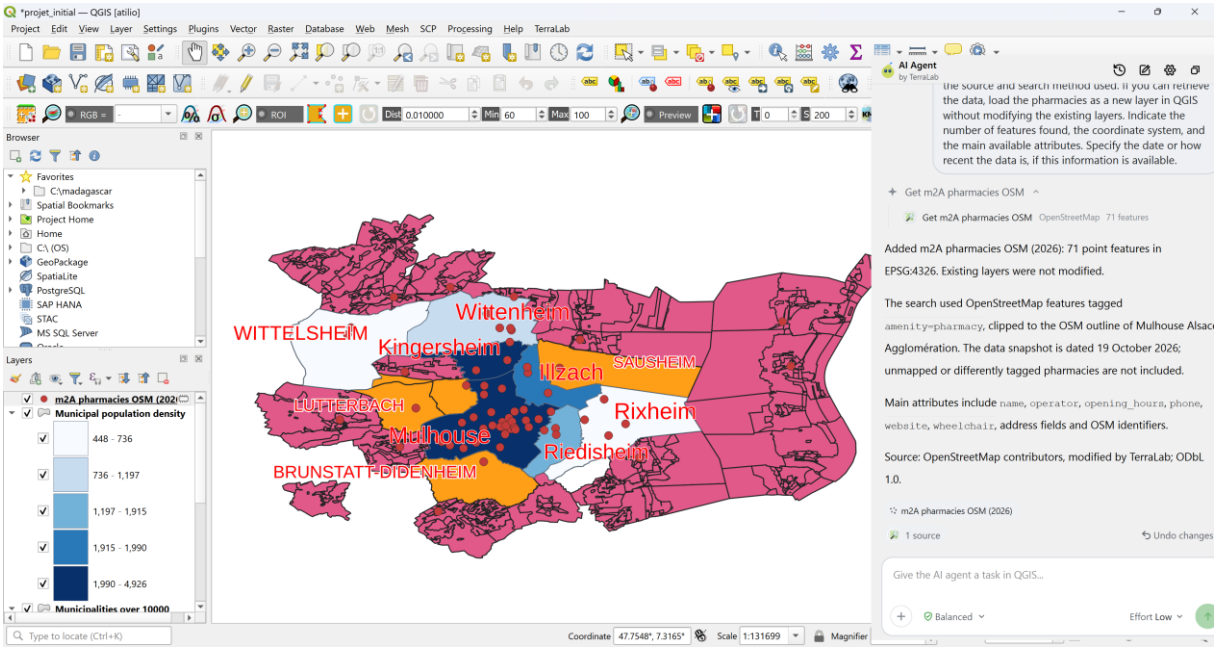
Added m2A pharmacies OSM (2026): 71 point features in EPSG:4326. Existing layers were not modified.

The search used OpenStreetMap features tagged `amenity=pharmacy`, clipped to the OSM outline of Mulhouse Alsace Agglomération. The data snapshot is dated 19 October 2026; unmapped or differently tagged pharmacies are not included.

Main attributes include `name`, `operator`, `opening_hours`, `phone`, `website`, `wheelchair`, address fields and OSM identifiers.

Source: OpenStreetMap contributors, modified by TerraLab; ODbL 1.0.

Result in QGIS :



Validation:

AI Agent did more than retrieve external data: it created a new point layer and loaded it directly into QGIS.

The resulting layer contains 71 pharmacy features in EPSG:4326, with a range of useful attributes, including names, operators, opening hours, contact details, addresses and OpenStreetMap identifiers. Existing project layers were not modified.

One point requires clarification: the execution report gives 19 October 2026 as the date of the OSM data snapshot, but does not specify whether this refers to the retrieval date, the extraction date, or the last modification date of the individual OSM objects. These are different types of information and should not be confused.

The report also identifies the data source and its licence (OpenStreetMap contributors, modified by TerraLab; ODbL 1.0).

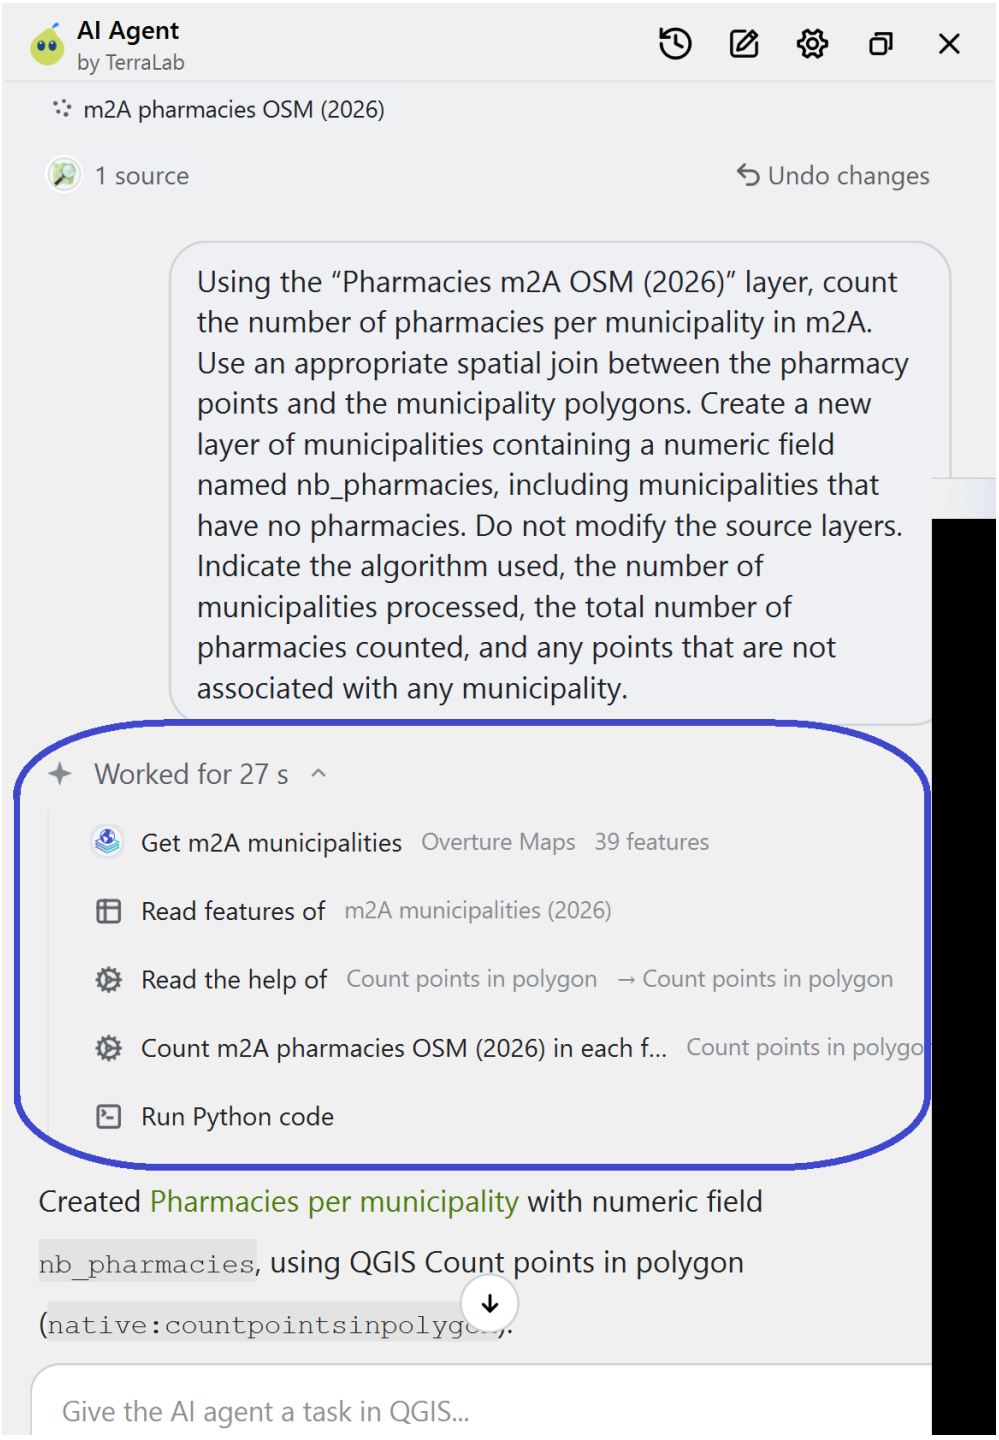
Test 13: Counting Pharmacies by Municipality

Objective: Count the number of pharmacies by municipality. This is a logical next step that will allow us to evaluate not only the search functionality but also the integration of external data into a geographic analysis.

Test prompt:

Using the "Pharmacies m2A OSM (2026)" layer, count the number of pharmacies per municipality in m2A. Use an appropriate spatial join between the pharmacy points and the municipality polygons. Create a new layer of municipalities containing a numeric field named nb_pharmacies, including municipalities that have no pharmacies. Do not modify the source layers. Indicate the algorithm used, the number of municipalities processed, the total number of pharmacies counted, and any points that are not associated with any municipality.

Execution :



Created Pharmacies per municipality with numeric field `nb_pharmacies`, using QGIS Count points in polygon (`native:countpointsinpolygon`).

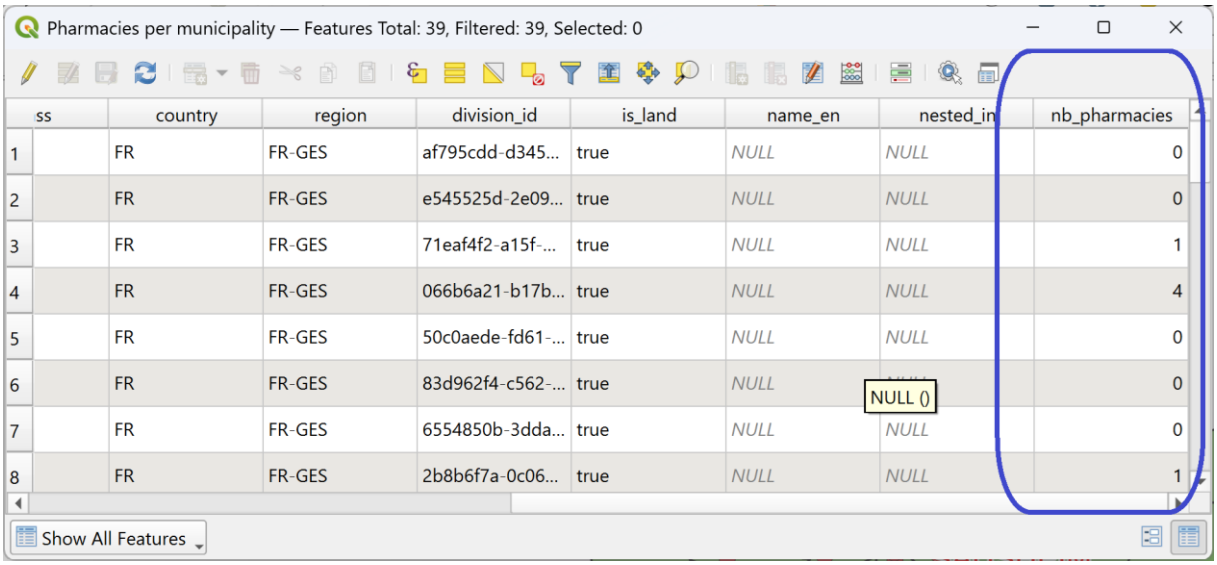
Municipalities processed: 39, including 16 with zero pharmacies.

Pharmacies counted: 71; unassociated points: 0; points in multiple municipalities: 0.

Existing source layers were not modified. The 11-feature municipality layer was filtered, so the count used newly loaded m2A municipality boundaries instead.

Source for boundaries: OpenStreetMap contributors, Overture Maps Foundation, geoBoundaries, Esri Community Maps contributors and LINZ, modified by TerraLab; ODbL 1.0.

Result in QGIS :



Validation:

AI Agent linked an external data search to a spatial analysis of municipalities.

A screenshot of the attribute table confirms that the resulting layer contains 39 municipalities and a numeric field named `nb_pharmacies`, with values including several zeros.

The execution report also confirms that all 71 pharmacies were counted, with no unassociated points and no points assigned to multiple municipalities.

An important detail is that the analysis used the newly loaded m2A municipality boundaries, rather than the original 11-feature municipality layer, which was filtered.

This test demonstrates the agent's ability to perform a sequence of operations: retrieve external data, create and load a new layer, and then run a QGIS geoprocessing algorithm to produce a spatial summary.

Test 14 — Retrieving Data from Another Source

Objective: To step outside of OpenStreetMap to verify whether AI Agent can query an institutional data source.

The official documentation mentions an IGN and DGFiP connector for French addresses, cadastral parcels, urban planning zones, building heights, and aerial photographs. Cadastral and urban planning data are limited to France and are released under the Open License 2.0.

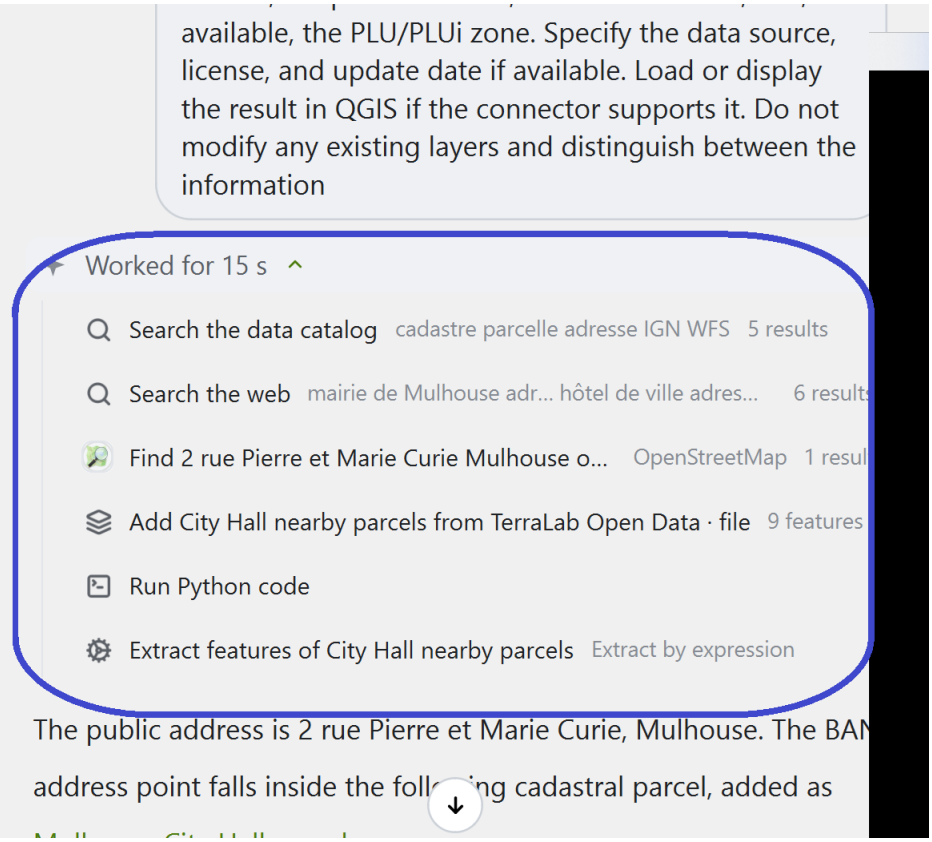
TerraLab also documents a tool dedicated to French cadastral parcels, capable of locating a parcel based on an address or a point, and then providing, among other things, its cadastral reference, municipality, section, number, area, and PLU/PLUi zone.

This does not necessarily mean that all cadastral parcels in a municipality can be downloaded in a single operation. This is precisely an interesting distinction to test.

Test prompt:

Search for the cadastral parcel corresponding to the Mulhouse City Hall in France, using its public address. Use the available French cadastral connector. Specify the complete cadastral reference, the municipality, the section, the parcel number, the cadastral area, and, if available, the PLU/PLUi zone. Specify the data source, license, and update date if available. Load or display the result in QGIS if the connector supports it. Do not modify any existing layers and distinguish between the information

Execution :



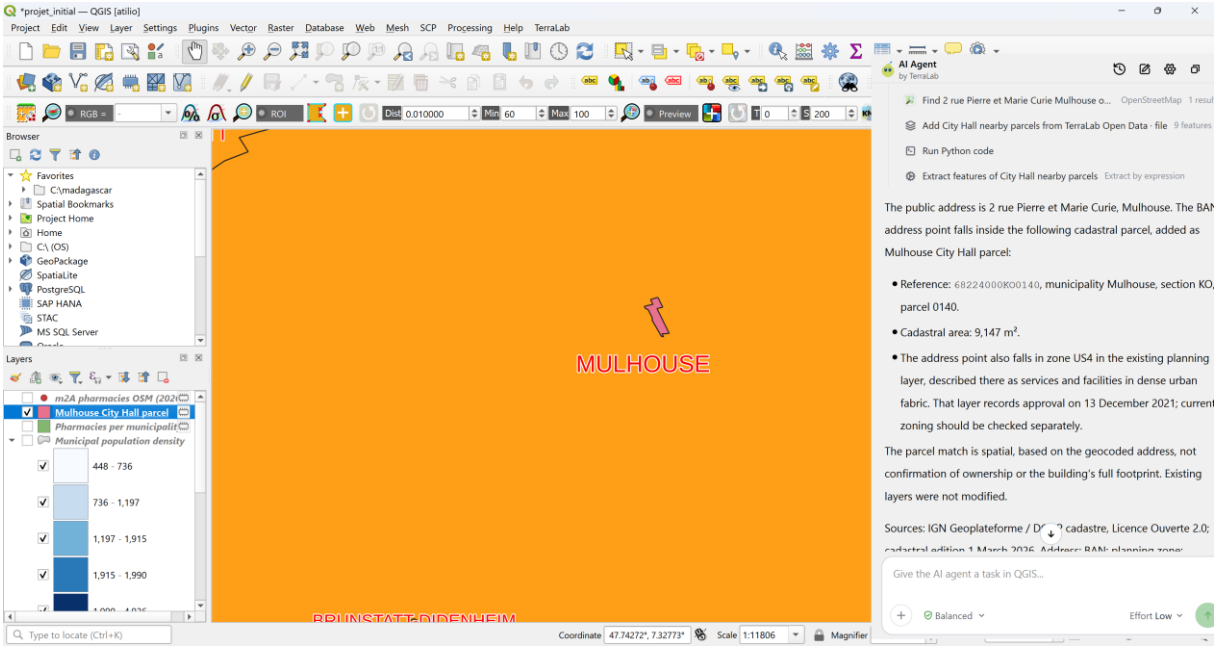
The public address is 2 rue Pierre et Marie Curie, Mulhouse. The BAN address point falls inside the following cadastral parcel, added as Mulhouse City Hall parcel:

- Reference: `68224000KO0140`, municipality Mulhouse, section KO, parcel 0140.
- Cadastral area: 9,147 m².
- The address point also falls in zone US4 in the existing planning layer, described there as services and facilities in dense urban fabric. That layer records approval on 13 December 2021; current zoning should be checked separately.

The parcel match is spatial, based on the geocoded address, not confirmation of ownership or the building's full footprint. Existing layers were not modified.

Sources: IGN Geoplateforme / DGFiP cadastre, Licence Ouverte 2.0; cadastral edition 1 March 2026. Address: BAN; planning zone: existing project layer, licence and update date not verified.

Result in QGIS:



Validation:

AI Agent successfully combined several geographic data sources and performed a sequence of spatial operations to identify the cadastral parcel associated with the town hall's address and cross-reference it with the existing planning layer.

What This Test Demonstrates

The agent did not merely search for textual information. According to its execution report, it performed several geographic processing steps:

Geocoded the town hall's address using the BAN.

Located the cadastral parcel containing the resulting address point.

Retrieved the cadastral reference and area.

Cross-referenced the address point with the existing planning layer and identified zone US4.

Loaded the results into QGIS.

This is a more advanced use case than a simple data query. It tests the integration of data sources, geocoding, cadastral information, spatial analysis and the creation of a GIS layer.

However, the results must be interpreted carefully. The cadastral match is based on the geocoded address point; it does not establish ownership or confirm the full footprint of the building. Similarly, the US4 zoning information comes from the existing project layer, whose current validity has not been independently verified. The layer records an approval date of 13 December 2021.

The test therefore demonstrates the agent's ability to combine geographic information and produce a spatial result, while also highlighting the need to distinguish retrieved data from information requiring further verification.

Test 15 — Verify the Geometry and Consistency of the Result

Objective: To verify what the agent actually loaded into QGIS, rather than asking them to confirm their own report.

Test prompt:

On the “Mulhouse City Hall Parcel” layer you just created, perform a consistency check without modifying the data.

1. Identify the geometry type, coordinate system, number of features, and available fields.
2. Check whether the layer contains a single parcel or multiple features.
3. If the data allows it, compare the cadastral area of 9,147 m² with the area calculated from the geometry, using a method appropriate for the coordinate system.
4. Check whether the geometry is valid.
5. Indicate whether the US4 zoning is stored within the parcel layer or comes from a separate layer.
6. If a cadastral or zoning source layer is available in the project, provide its name and any accessible provenance information.

Do not correct any geometry, create any fields, or modify any layers. Distinguish between checks that were actually performed and information that was simply carried over from the previous result.

Execution :

Checks performed on Mulhouse City Hall parcel:

- One MultiPolygon feature, in EPSG:4326. Its geometry passes the GEOS validity check.
- Fields: `fid`, `id`, `gid`, `numero`, `feuille`, `section`, `code_dep`, `nom_com`, `code_com`, `com_abs`, `code_arr`, `idu`, `contenance`, `code_insee`.

- The stored cadastral area is 9,147 m². The geometry measures 9,148.85 m² on the WGS 84 ellipsoid, 1.85 m² more.
- US4 is not stored in the parcel layer. The previous address-point check found it in the separate zone_urba_mulhouse layer, sourced from the project GeoPackage; its accessible metadata has no title, abstract or history.

City Hall nearby parcels is the cadastral source layer in the project. Its IGN Geoplateforme origin and Licence Ouverte 2.0 are carried over from the earlier retrieval, not reverified here. No data or layers were modified.

Validation :

This test provides an important consistency check. AI Agent did not simply reuse the cadastral area: it calculated the area of the parcel geometry on the WGS 84 ellipsoid and compared it with the cadastral value.

Evaluation of Test 15 — Consistency Check

Check	Result	Evaluation
Number of features	1	Consistent with a single parcel
Geometry	MultiPolygon	Consistent with a cadastral parcel
CRS	EPSG:4326	Geographic coordinates; direct planar area calculations would not be appropriate
Geometric validity	Valid according to GEOS	Validity check reported as passed
Cadastral area	9,147 m²	Value retrieved from cadastral data
Geodetic area	9,148.85 m²	Independently calculated from the geometry
Difference	+1.85 m², approximately +0.02%	Very small difference
Data modification	None	Source data and project layers preserved

The Agent Correctly Distinguishes Between Layers

Another positive point is that the US4 zoning designation does not appear in the parcel's attributes. The agent identifies it as information held in a separate layer, `zone_urba_mulhouse`, from the project's GeoPackage.

This distinction is important: cadastral information and urban planning information come from different layers and should not be presented as if they were attributes of the same dataset.

What Remains to Be Verified

- The execution report is consistent, but some aspects would benefit from independent verification:
- Whether the cadastral source layer contains the correct parcel and whether its geometry was preserved during loading.
- Whether the reported GEOS validity check can be independently reproduced on the parcel geometry.
- Whether the entire parcel is covered by the US4 zone, rather than only the address point being located within that zone.

The last point is particularly important: the previous test established that the geocoded address point falls within US4, but this execution does not report a full parcel-versus-zone intersection or coverage calculation.

These reservations concern the independent verification and interpretation of the results, not an error identified in the agent's reported calculations.

Test 16 — Verify Zoning Through Spatial Analysis

Objective: To test whether the agent can quantify the overlap between the parcel and the zoning designation, rather than simply stating that it covers the entire parcel.

Test prompt:

Using the "Mulhouse City Hall Parcel" and "PLU Zoning Near City Hall" layers, quantitatively verify the overlap of the parcel with zone US4.

1. *Identify the geometries and coordinate systems of the two layers.*
2. *Verify whether the parcel intersects one or more zoning features.*
3. *Calculate the area of the parcel covered by US4 and any area that may lie outside this zone, using an appropriate method.*
4. *Provide the percentage of the parcel covered by US4.*
5. *If multiple zoning features intersect the parcel, verify that their areas are not added together with duplicate overlaps.*
6. *Indicate the QGIS algorithms or expressions used and the limitations of the check.*

Do not modify any existing layers or create any new fields. If a required layer or geometry is not available, indicate this without inventing a result.

Execution :

- Mulhouse City Hall parcel and the zoning layer zone_urba_mulhouse are both MultiPolygon layers in EPSG:4326. The latter was used in place of the named "PLU Zoning Near City Hall" layer.
- The parcel intersects one zoning feature, US4. Both geometries passed the validity check.
- In Lambert-93 (EPSG:2154), the calculated parcel area is 9,135.73 m²: all 9,135.73 m², or 100%, is covered by US4; 0 m² lies outside it. The US4 intersection was unioned before measurement, preventing double counting.

The check used PyQGIS intersection, unary union, difference and planar area operations; no layers or fields were changed. The planar result differs from the previously reported ellipsoidal area and the cadastral area of 9,147 m². It reflects the geometry and zoning data available in the project, not a confirmation of current legal zoning.

Validation:

The agent moves from a qualitative zoning check to a quantitative spatial analysis, calculating both the area covered by the zoning and the area outside it.

What This Test Demonstrates

Three elements are worth noting in the evaluation of AI Agent QGIS:

- **The agent can chain together multiple geometric operations.** It uses intersection to calculate the common area, unary union to combine the intersecting geometries and prevent double counting, and difference to determine the area outside the zoning. This is more robust than a simple visual inspection.
- **The agent provides a complete area assessment.** In Lambert-93 (EPSG:2154), the parcel area is calculated as 9,135.73 m². The intersection with US4 covers the entire parcel: 9,135.73 m², or 100%, with no area outside the zone. The result is arithmetically consistent.
- **The area calculation must be interpreted in context.** The 9,135.73 m² planar area differs from the 9,148.85 m² ellipsoidal area calculated in the previous test and from the cadastral area of 9,147 m². These values result from different measurement methods and should not be treated as interchangeable.

The agent also clarifies the scope of its conclusion. The result applies to the geometries available in the project. It does not establish the current legal validity of the zoning designation or replace consultation of the applicable planning regulations, particularly when assessing a building or development project.

Test 17 — Testing a Cadastral Search at the Municipal Level

Objective: To test a cadastral search at the municipal level. So far, we have tested searching for a parcel based on an address, followed by checks on that parcel. One important practical question remains: Can the user retrieve a set of parcels, or only identify a single parcel?

Test prompt:

I want to retrieve the cadastral parcels for the municipality of Mulhouse (INSEE code 68224) in QGIS.

Search the available connectors and data sources to see if any provide access to retrieve all cadastral parcels for the municipality, rather than just a single parcel based on an address or a point.

Specify the name of the connector or service, the access method, the geographic coverage, and the license terms.

Specify whether the service actually allows you to download all parcels in the municipality in a single operation, or if you need to use queries by right-of-way, by lot, or by identifier.

If a complete download is possible, indicate the approximate number of expected features and any limitations before starting the operation.

If a complete download is not possible with the available tools, explain the limitation and suggest a realistic alternative method.

For this test, do not download all the parcels yet, do not create any layers, and do not modify the project. Clearly distinguish between the capabilities confirmed by the connector and those that are merely assumed.

Execution :

Cadastre Etalab, available through the IGN France connector, provides a whole-commune GeoJSON file by INSEE code. Its catalog confirms French coverage and a single-file download for commune 68224; individual parcel or spatial queries are not required. The licence is Licence Ouverte 2.0, with Etalab and DGFiP attribution.

A count-only query to the separate IGN Geoplateforme cadastral WFS returned 22,815 parcels for `code_insee = 68224`. This is a useful expectation, not a verified count of the Etalab file: editions may differ.

The WFS is an alternative, but it limits responses to 5,000 features per request, so a complete raw WFS retrieval requires paging. The whole-commune Etalab download itself was not tested.

No parcels were downloaded, no layers were created, and the project was not modified..

Validation:

AI Agent identified two ways to access municipal cadastral data, each with different retrieval methods and constraints.

What the agent did well:

- It identified two complementary access methods: a complete municipal GeoJSON download through Cadastre Etalab and queries to the IGN cadastral WFS.
- It queried the WFS to obtain a count of 22,815 parcels without downloading the full dataset.
- It correctly identified the WFS limit of 5,000 features per request and the need to handle pagination for a complete retrieval.
- It did not present the WFS count as a verified count of the Etalab GeoJSON file, acknowledging that the two sources may differ in date or content.
- It respected the instruction not to download parcels, create a layer or modify the project.

One important distinction is that the Etalab download was identified in the catalogue but was not actually tested. The test therefore evaluates the agent's ability to identify and compare data access methods, rather than its ability to retrieve and load the complete municipal dataset.

The distinction between the WFS count and the expected number of features in the GeoJSON file is important: the two sources may differ in terms of date, content or distribution method.

.

Test 18 — Compare the Two Sources Without Loading the Cadastral Parcels

Objective: Compare the two sources without loading the cadastral parcels. This test evaluates the agent's ability to research and assess their sources before initiating a process. It also provides an opportunity to verify whether the agent can recognize the limitations of their catalog and avoid confusing two services that appear to provide the same data.

Test prompt:

Compare the two sources of cadastral parcels identified for the municipality of Mulhouse (INSEE code 68224):

- the municipal GeoJSON file from the Cadastre Etalab via the IGN France connector;
- the WFS service from IGN Géoplateforme, filtered by code_insee = 68224.

Without downloading the parcels or modifying the project:

1. Retrieve, if the services allow it, the available metadata for each source: data date or year, update date, coordinate system, attribute fields, access terms, and license.

2. Verify whether the number of 22,815 parcels is confirmed by a response from the WFS service and specify the method used to obtain this figure.

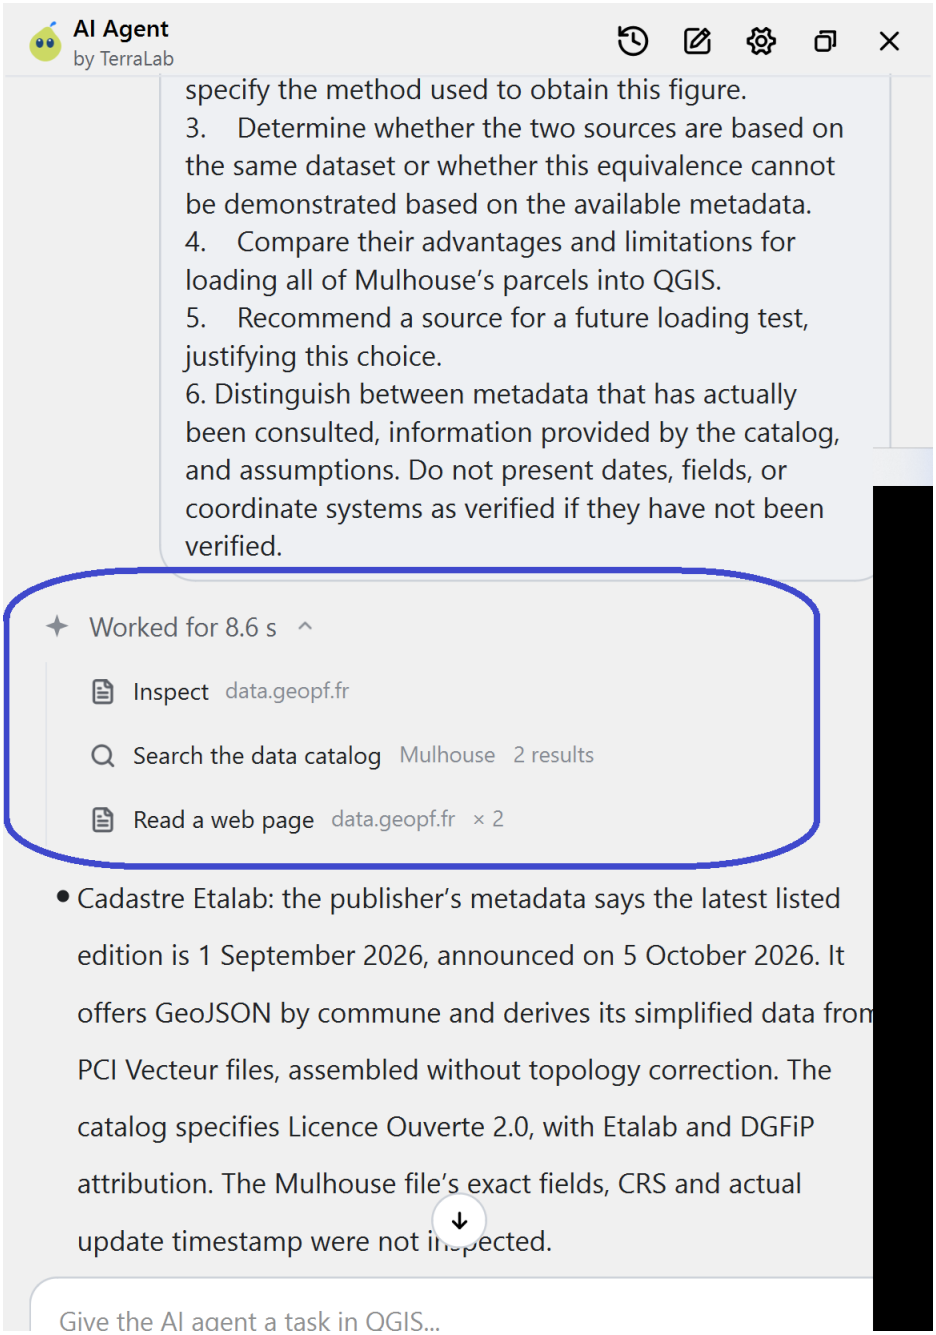
3. Determine whether the two sources are based on the same dataset or whether this equivalence cannot be demonstrated based on the available metadata.

4. Compare their advantages and limitations for loading all of Mulhouse's parcels into QGIS.

5. Recommend a source for a future loading test, justifying this choice.

6. Distinguish between metadata that has actually been consulted, information provided by the catalog, and assumptions. Do not present dates, fields, or coordinate systems as verified if they have not been verified.

Execution :



Cadastre Etalab: the publisher's metadata says the latest listed edition is 1 September 2026, announced on 5 October 2026. It offers GeoJSON by commune and derives its simplified data from PCI Vecteur files, assembled without topology correction. The catalog specifies Licence Ouverte 2.0, with Etalab and DGFiP attribution. The Mulhouse file's exact fields, CRS and actual update timestamp were not inspected.

IGN Géoplateforme WFS: a fresh count-only service query filtered by `code_insee = 68224` returned 22,815 matches. No parcel geometries were downloaded. The service's schema lists `gid`, `numero`, `feuille`, `section`, `code_dep`, `nom_com`, `code_com`, `com_abs`, `code_arr`, `idu`, `contenance`, `code_insee` and `geom`. The catalog reports edition 1 March 2026, Licence Ouverte

2.0 and native Lambert-93 (EPSG:2154); the service's current edition and default CRS were not independently verified.

The catalog describes both as derived from DGFiP cadastral parcels, but identical records, geometry and edition have not been demonstrated. The publisher describes Etalab as a reworked, simplified distribution.

For a future test, I recommend Etalab's commune GeoJSON: one municipality-wide file avoids WFS paging. The catalog reports a 5,000-feature WFS response limit, so loading all 22,815 through raw WFS requires paging and a count check. The Etalab file's final count must also be checked after loading.

No parcels were downloaded and the project was not modified.

Validation :

This test confirms that the agent can compare two cadastral data sources while taking into account their origin, publication dates, metadata and access methods. It correctly avoids concluding that two datasets derived from the same DGFiP cadastral data are identical at parcel level.

Choosing the Etalab GeoJSON file makes sense for the next test: a single municipality-wide file should simplify loading into QGIS, whereas retrieving all 22,815 WFS features requires managing multiple pages.

However, several points still need to be checked during the download:

- **Completeness:** Does the downloaded file contain all the expected parcels for Mulhouse? The WFS count of 22,815 provides a useful comparison, but does not establish the exact expected count in the Etalab file.
- **Structure:** What are the actual CRS, available fields, geometry types and any missing values in the downloaded file?
- **Data currency:** Does the actual file reflect the latest cadastral edition? The more recent edition date announced for Etalab does not, by itself, prove that every parcel is more up to date than in the WFS.

Differences between the two datasets may also result from their respective data models and distribution methods. These aspects will need to be assessed using the downloaded data.

Test 19 — Load the Etalab Municipal Cadastral Database

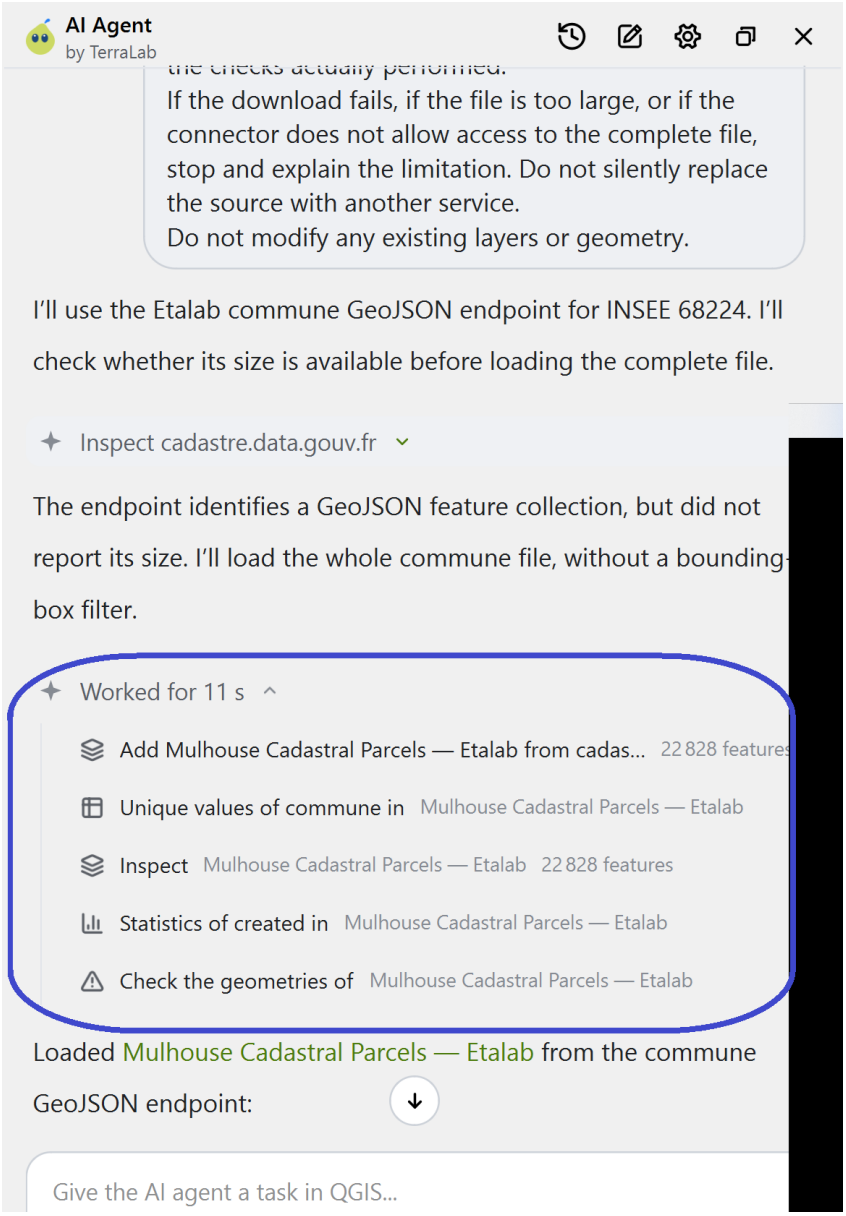
Objective: To determine whether the capacity listed in the catalog actually results in a complete and verifiable municipal dataset in QGIS

Test prompt:

Download and load into QGIS the Etalab municipal GeoJSON file for Mulhouse (INSEE code 68224) using the IGN France connector identified earlier.
Before you begin, specify the URL or access method used and, if possible, verify the file size.
After loading:

Create a new layer named “Mulhouse Cadastral Parcels — Etalab” without modifying any existing layers.
Specify the total number of features, the geometry type, the CRS declared by the layer, and the attribute fields.
Check for the presence of the INSEE code 68224 in the features and flag any missing or differing values.
Verify geometric validity on a sample or on all features, if reasonable to do so.
Compare the number of features obtained with the 22,815 parcels reported by the WFS, without assuming that the two numbers must necessarily be identical.
Indicate the source, year of creation, and license, distinguishing between the producer’s metadata and the checks actually performed.
If the download fails, if the file is too large, or if the connector does not allow access to the complete file, stop and explain the limitation. Do not silently replace the source with another service.
Do not modify any existing layers or geometry.

Execution :



I'll use the Etalab commune GeoJSON endpoint for INSEE 68224. I'll check whether its size is available before loading the complete file.

The endpoint identifies a GeoJSON feature collection, but did not report its size. I'll load the whole commune file, without a bounding-box filter.

Loaded Mulhouse Cadastral Parcels — Etalab from the commune GeoJSON endpoint:

cadastre.data.gouv.fr

The original download size was not reported; the imported GeoPackage is 8,179,712 bytes.

The layer has 22,828 polygon features in EPSG:4326. All 22,828 have `commune = 68224`; none are missing or different. All passed the geometry validity check.

Fields: `fid`, `id`, `commune`, `prefixe`, `section`, `numero`, `contenance`, `arpente`, `created`, `updated`.

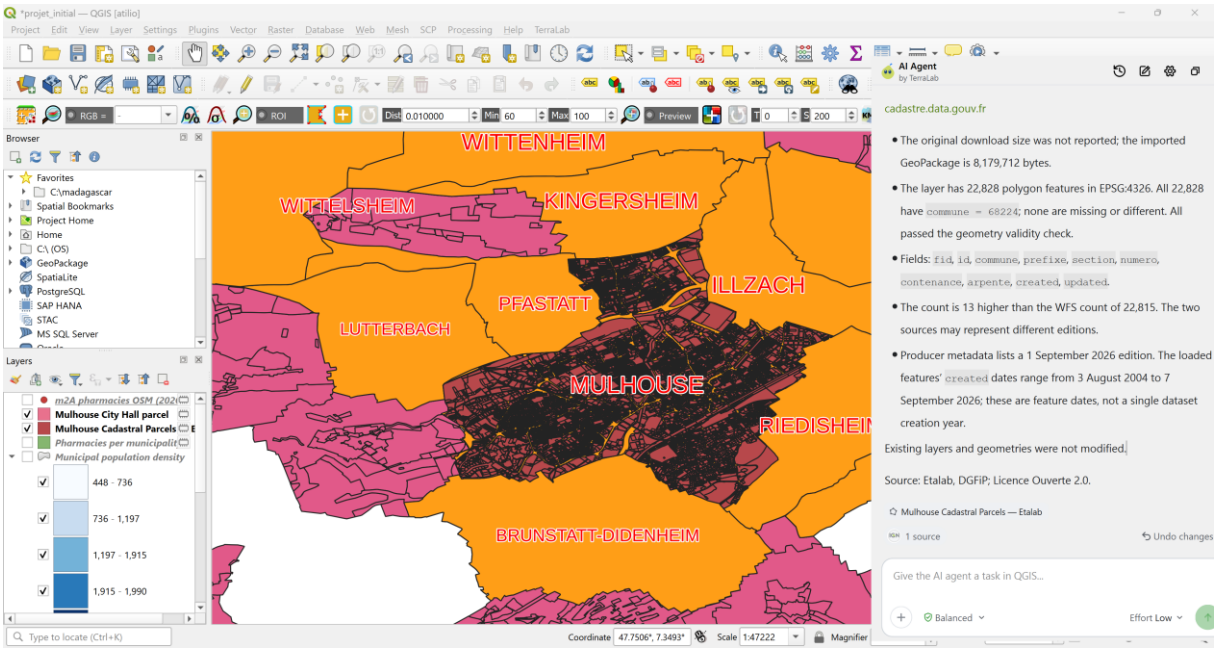
The count is 13 higher than the WFS count of 22,815. The two sources may represent different editions.

Producer metadata lists a 1 September 2026 edition. The loaded features' `created` dates range from 3 August 2004 to 7 September 2026; these are feature dates, not a single dataset creation year.

Existing layers and geometries were not modified.

Source: Etalab, DGFiP; Licence Ouverte 2.0.

Result in QGIS:



Validation :

AI Agent successfully loaded a complete municipal cadastral dataset containing 22,828 polygon features and performed several checks on the resulting layer. We have now moved beyond simply searching for a parcel to testing the creation and inspection of a full-fledged GIS dataset.

1. A Real Change in Scale

This test shows that the agent can:

Retrieve a municipal file from an external source.

Import a GeoJSON file into a local GeoPackage.

Load the resulting layer into QGIS.

Examine its attributes and coordinate reference system.

Check the municipality code and missing values.

Run a geometry validity check on the loaded features.

This result is more conclusive than the previous tests, as it involves a complete municipal dataset rather than an individual parcel.

2. The Discrepancy of 13 Parcels Warrants Further Investigation

The two sources provide different counts:

Etalab: 22,828 parcels.

IGN Geoplateforme WFS: 22,815 parcels.

A difference in publication dates or cadastral editions is plausible, but has not been demonstrated. Other explanations are possible, including differences in filtering, data scope, data model or counting methods.

The 13 additional features should therefore not be considered an error, nor evidence of a more recent update, until the two datasets have been compared.

3. What the Data Checks Actually Confirm

The execution report states that all 22,828 geometries passed the geometry validity check. It also confirms that all features have the municipality code 68224, with no missing or different values.

These are useful quality checks, but they do not independently establish that the dataset contains every cadastral parcel belonging to Mulhouse. Completeness would require comparison with an authoritative reference or an independent count for the same edition and scope.

The report also distinguishes the producer's edition date (1 September 2026) from the individual feature creation dates, which range from 3 August 2004 to 7 September 2026. These dates should not be confused with a single dataset creation date.

Test 20 — Comprehensive Check of the Etalab Layer

Objective: To evaluate a key skill for professional use: the agent's ability to review a large dataset and present a quantified quality assessment, without confusing a sample check with a comprehensive check.

Test prompt:

On the "Mulhouse Cadastral Parcels — Etalab" layer, perform a quality check on all 22,828 features without modifying the layer or its geometries.

Verify the total number of features and the number of null, empty, or invalid geometries.

If possible, use the QGIS geometric validity check algorithm or an equivalent PyQGIS check on all features. Indicate the method used.

Check for null or empty values in the section, number, area, and municipality fields.

Check for any duplicates based on the combination of the municipality, prefix, section, and number fields. Explain the limitations of this key if necessary.

Verify the minimum and maximum values and the number of null values for area. Do not assume that a null area value is necessarily an error.

Provide a summary of the checks with numerical data and explicitly indicate which checks could not be performed.

Do not correct any geometry, delete any entities, create any fields, or modify the project. If the full processing is likely to be too costly, explain this before initiating a resource-intensive operation.

Execution :

Features checked: 22,828. Null geometries: 0; empty: 0; invalid: 0, using the full-layer geometry validity check.

Null or empty values: `section`, `numero` and `commune` each have 0. `contenance` has 1 null and 0 empty values.

Non-null cadastral areas range from 1 to 278,640 m². The null value is not necessarily an error.

Duplicate keys using `commune`, `prefixe`, `section` and `numero`: 0. This checks exact attribute matches, not duplicate geometries or historical parcel changes.

All requested checks were completed. No data, geometry or project state was modified.

Validation :

AI Agent performed a comprehensive quality check on all 22,828 features without modifying the data. This is an important result when assessing its usefulness in professional geomatics work.

What We Learned

- **Full-layer geometry validation.** The agent checked all 22,828 geometries, reporting no null, empty or invalid geometries. This confirms its ability to perform a complete geometry validity check rather than relying on a sample.
- **Distinguishing missing values from zero.** The agent identified one null value in the `contenance` field, with no empty values. It correctly noted that a missing cadastral area is not necessarily an error. This distinction matters: a null value represents missing information, whereas a numerical value of zero is an explicitly recorded value with a different meaning.
- **Understanding the limits of uniqueness checks.** No duplicate combinations were found for `commune`, `prefixe`, `section` and `numero`. This confirms the uniqueness of these attribute combinations within the loaded layer. However, it does not guarantee the accuracy of cadastral references, the absence of duplicate geometries or the completeness of the dataset.

All requested checks were completed without modifying the source data, geometries or project state.

For independent validation, we compared the results with those of a separate QGIS check, and the results were confirmed.

Test 21 — Generate a Statistical Summary of the Cadastral Register

Objective: We have verified the structure and geometric quality. The next step is to test the agent's ability to summarize the data without altering its content.

Test prompt:

Using the "Mulhouse Cadastral Parcels — Etalab" layer, generate a descriptive statistical summary without modifying the layer or creating any fields.

Provide the total number of parcels and the number of parcels per cadastral section.

Calculate the count, mean, median, minimum, maximum, and sum of the reported areas.

Indicate separately the number of zero areas and the number of values equal to zero.

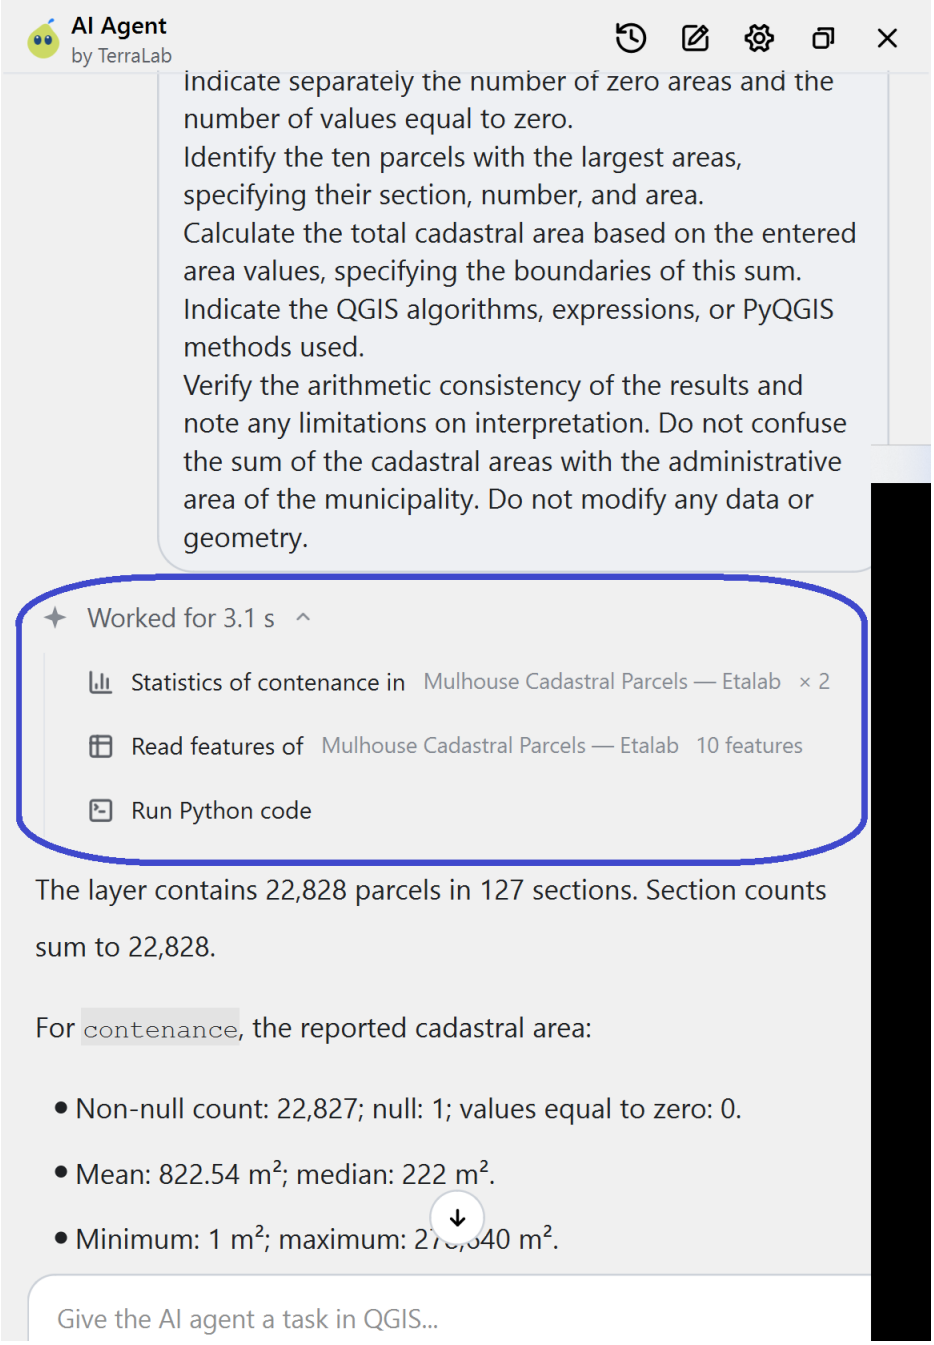
Identify the ten parcels with the largest areas, specifying their section, number, and area.

Calculate the total cadastral area based on the entered area values, specifying the boundaries of this sum.

Indicate the QGIS algorithms, expressions, or PyQGIS methods used.

Verify the arithmetic consistency of the results and note any limitations on interpretation. Do not confuse the sum of the cadastral areas with the administrative area of the municipality. Do not modify any data or geometry.

Execution :



The layer contains 22,828 parcels in 127 sections. Section counts sum to 22,828.

For `contenance`, the reported cadastral area:

Non-null count: 22,827; null: 1; values equal to zero: 0.

Mean: 822.54 m²; median: 222 m².

Minimum: 1 m²; maximum: 278,640 m².

Sum: 18,776,120 m², or 18.77612 km².

The sum includes every non-null reported area in the layer, including zero if present; it excludes the one null. It is not the municipality's administrative area and is not a geometric union. The independently summed values and section counts agree with the statistics.

Ten largest reported areas, all with prefix 000:

LP 32: 278,640 m²

ND 3: 215,665 m²

NC 63: 213,155 m²

LT 3: 198,625 m²

NE 79: 189,972 m²

LA 77: 125,613 m²

NO 321: 122,301 m²

MZ 351: 122,285 m²

MY 30: 110,916 m²

KB 256: 107,351 m²

Parcel counts by section:

EM 31; EN 52; EO 61; EP 113; ER 98; ES 187; ET 146; EV 120; EW 139; EX 32; EY 63; EZ 175

HA 24; HB 31; HC 291; HD 113; HE 74; HI 27; HK 15; HL 80; HM 69; HN 123; HO 131; HP 73; HR 167; HS 150; HT 218; HV 256; HW 172; HX 521; HY 235; HZ 435

IA 85; IB 78; IC 167; ID 101; IE 165; IH 555; IK 128; IL 153; IM 178; IN 62; IP 117; IR 383; IS 160; IT 165; IV 397; IW 609; IX 349; IY 466; IZ 233

KA 271; KB 55; KC 60; KD 199; KE 220; KH 122; KI 85; KK 92; KL 168; KM 297; KN 81; KO 110; KP 222; KR 179; KS 187; KT 93; KV 123; KW 177; KX 115; KY 176; KZ 258

LA 25; LB 330; LC 222; LD 75; LE 178; LH 368; LI 356; LK 340; LL 314; LM 212; LN 267; LO 144; LP 26; LR 101; LS 76; LT 16; LV 90; LW 244; LX 101; LY 176; LZ 191

MA 266; MB 230; MC 258; MD 108; ME 389; MH 181; MI 191; MK 213; ML 171; MM 211; MN 339; MO 120; MP 238; MR 235; MS 193; MT 185; MV 236; MW 230; MX 180; MY 25; MZ 281

NA 108; NB 273; NC 75; ND 10; NE 90; NH 259; NI 274; NK 101; NL 196; NM 290; NN 245; NO 278; NP 34

QGIS field statistics and descending attribute sorting supplied the distribution and largest parcels; a read-only PyQGIS pass counted sections, nulls and zeros and independently summed areas. No data or geometry was modified.

Validation :

The agent produced a comprehensive statistical report on the 22,828 parcels, including a breakdown by cadastral section, descriptive statistics for parcel areas, and a ranking of the ten largest parcels. It also explained the calculation methods and the limitations of interpretation.

The reported average is consistent with the sum and the number of non-null values. The unit conversion is also correct: 18,776,120 m² corresponds to 1,877.612 hectares, or 18.77612 km².

The independent checks reported in the execution confirm that the section counts sum to 22,828 and that the independently calculated sum of cadastral areas agrees with the reported statistics.

The agent also correctly points out that the sum of cadastral areas is not necessarily equal to the municipality's administrative area. These are different measurements: the former is the sum of recorded parcel areas, not the area of the municipality or the result of a geometric union of parcel boundaries.

This test demonstrates the agent's ability to produce and cross-check a substantial statistical summary from a large cadastral dataset.

Test 22 — Extract Large Parcels and Verify the Results

Objective: This test is designed to assess whether the user can convert statistics into a usable GIS selection without modifying the original cadastral layer.

Test prompt:

Using the "Mulhouse Cadastral Parcels — Etalab" layer, identify the parcels with a cadastral area greater than or equal to 10,000 m² (1 hectare).

Indicate the number of parcels that meet this criterion.

Calculate the sum of the areas of these parcels, their average area, and their maximum area.

List the ten largest parcels in this selection, including municipality, prefix, section, number, and area.

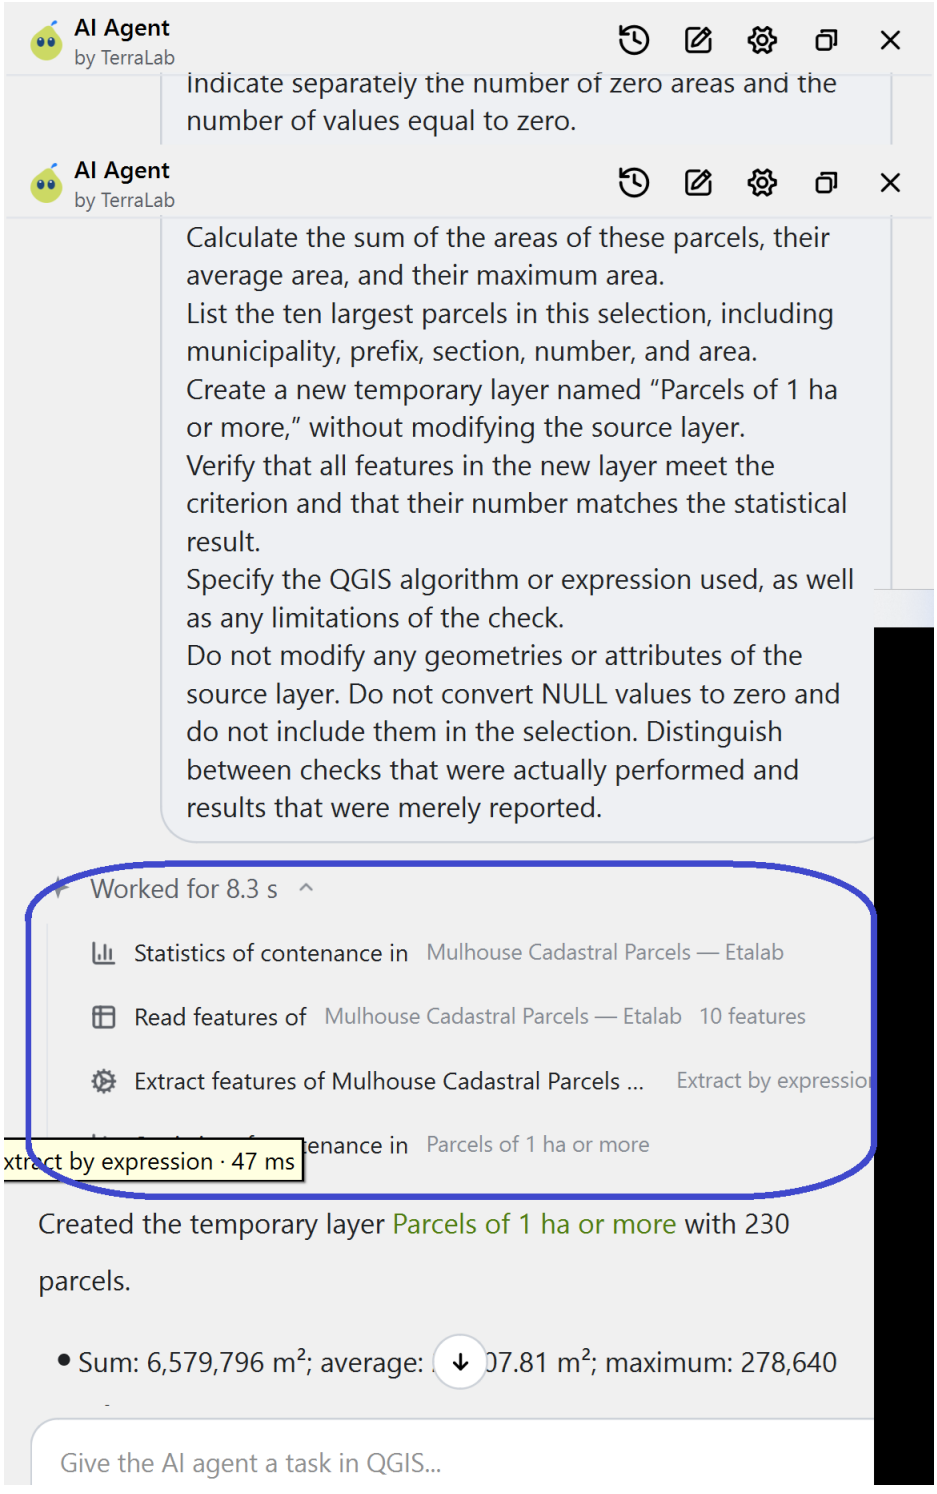
Create a new temporary layer named "Parcels of 1 ha or more," without modifying the source layer.

Verify that all features in the new layer meet the criterion and that their number matches the statistical result.

Specify the QGIS algorithm or expression used, as well as any limitations of the check.

Do not modify any geometries or attributes of the source layer. Do not convert NULL values to zero and do not include them in the selection. Distinguish between checks that were actually performed and results that were merely reported.

Execution :



Created the temporary layer Parcels of 1 ha or more with 230 parcels.

Sum: 6,579,796 m²; average: 28,607.81 m²; maximum: 278,640 m².

The ten largest are all in Mulhouse (68224), prefix 000:

LP 32: 278,640 m²

ND 3: 215,665 m²

NC 63: 213,155 m²

LT 3: 198,625 m²

NE 79: 189,972 m²

LA 77: 125,613 m²

NO 321: 122,301 m²

MZ 351: 122,285 m²

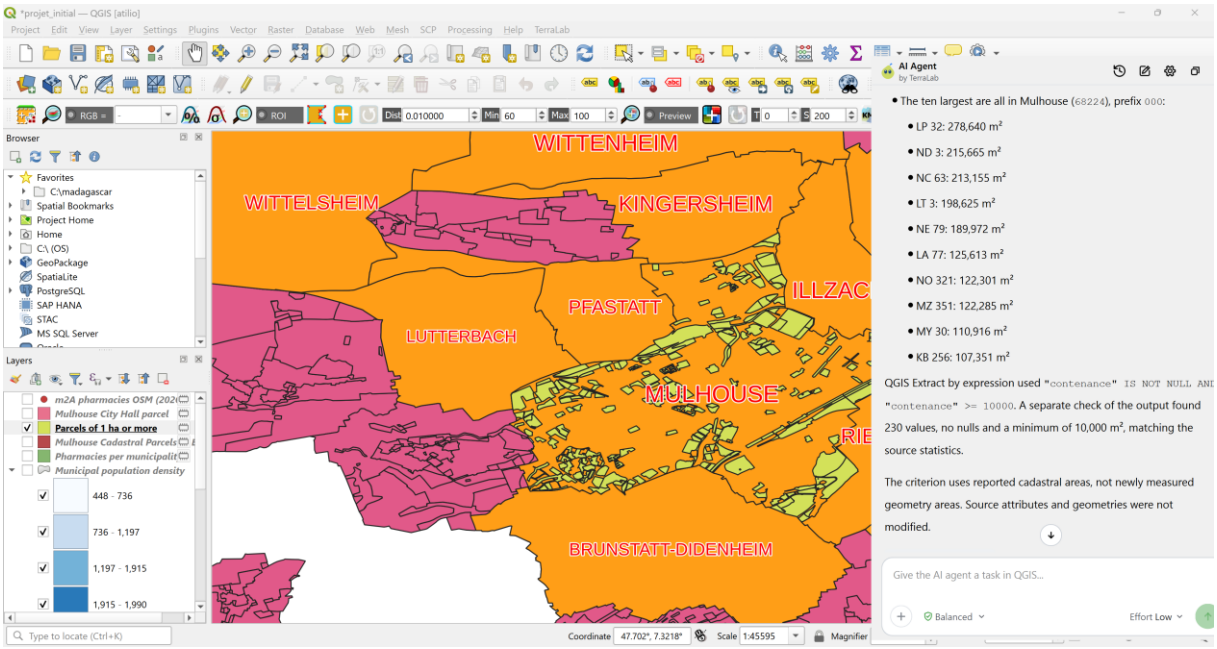
MY 30: 110,916 m²

KB 256: 107,351 m²

QGIS Extract by expression used `"contenance" IS NOT NULL AND "contenance" >= 10000`. A separate check of the output found 230 values, no nulls and a minimum of 10,000 m², matching the source statistics.

The criterion uses reported cadastral areas, not newly measured geometry areas. Source attributes and geometries were not modified.

Result in QGIS:



Validation:

AI Agent performed an attribute-based extraction, calculated statistics on the resulting layer, and checked the consistency of the output without modifying the source data.

The reported average is arithmetically consistent with the sum of the cadastral areas and the 230 selected features.

The ten largest parcels are the same as those identified in Test 20. This is consistent, since all ten exceed the 10,000 m² threshold.

An important methodological detail is that the extraction uses the recorded cadastral area (`contenance`), through the expression `"contenance" IS NOT NULL AND "contenance" >= 10000`. It does not use newly calculated geometric areas.

What This Test Contributes to the Overall Evaluation

At this stage, the tests show that the agent can:

- Search for and retrieve geographic and cadastral data.
- Load datasets into QGIS.
- Calculate statistics and check geometry validity.
- Perform attribute-based extractions.
- Create new layers and verify properties of the resulting data.
- Follow instructions not to modify source data.

The output layer was independently checked: it contains 230 features, no null values in `contenance`, and a minimum value of 10,000 m². The reported statistics and the results of these checks are consistent.

Test 23 — Combining an Attribute-Based Criterion and a Spatial Criterion

Objective: To verify whether the agent can combine two types of selection: identify parcels of at least one hectare that intersect a given urban planning zone.

Test prompt:

Using the cadastral and zoning layers already present in the project, find the cadastral parcels in Mulhouse with an area greater than or equal to 10,000 m² that intersect zone US4.

Verify that the two required layers are available and identify their names, CRS, and relevant fields.

Apply the attribute filter to the cadastral area, excluding NULL values.

Then apply the spatial filter to check for intersection with zone US4.

Create a new temporary layer named "Large Parcels Intersecting US4," without modifying the source layers.

Report the number of parcels obtained, their total cadastral area, and the ten largest parcels.

Specify whether the spatial criterion means simply that a parcel touches zone US4 or whether its entire geometry is located within that zone.

Verifies that each feature in the final layer satisfies both criteria.

If the available zoning layer covers only the city hall district and not the entire city of Mulhouse, notes this limitation and does not present the result as a complete inventory of parcels in the US4 zone throughout the municipality. Does not modify any source layers or invent any results if the necessary data arer not available.

Execution :

Created Large Parcels Intersecting US4 with 2 parcels, totaling 36,714 m² of reported cadastral area. Both were independently checked against the area and intersection criteria:

Mulhouse `68224`, prefix `000`, section LL, parcel 155: 25,498 m²

Mulhouse `68224`, prefix `000`, section ME, parcel 67: 11,216 m²

The inputs were Mulhouse Cadastral Parcels — Etalab (`contenance`, `commune`, `section`, `numero`) and zone_urba_mulhouse (`libelle`, `code_insee`), both in EPSG:4326. QGIS Extract by expression excluded null areas and kept those ≥10,000 m²; Extract by location then tested intersection with four Mulhouse US4 features.

Intersection includes boundary touches and does not mean the entire parcel is inside US4. The zoning layer extends beyond the city-hall district, but its completeness and currency across Mulhouse were not verified, so this is not a confirmed citywide inventory. No source layer was modified.

Validation :

This test combines two successive spatial selection operations and produces a non-empty result: two parcels with a combined reported cadastral area of 36,714 m².

The agent first selected parcels with a recorded cadastral area of at least 10,000 m², excluding null values. It then selected those intersecting the four US4 features in the available zoning layer.

The two resulting parcels were independently checked:

Section LL, parcel 155: 25,498 m².

Section ME, parcel 67: 11,216 m².

What This Test Demonstrates

The agent correctly combined an attribute-based extraction with a spatial selection. It also identified two important limitations when interpreting the result.

First, intersection does not mean complete inclusion. A parcel is selected if its geometry touches or overlaps the zoning geometry, even partially. This must be distinguished from a test that would require the entire parcel to lie within US4.

Second, the available zoning layer does not necessarily cover the whole municipality of Mulhouse. The two selected parcels therefore cannot be presented as an exhaustive inventory of all large parcels located in US4 across the city.

The result is valid for the available datasets and the specified selection criteria, but its geographical scope must be kept in mind.

Test 24 — Distinguishing Between Data Limitations and Geographic Results

Objective: Ask the agent to explain how to assess the reliability and geographical scope of the intersection results obtained in Test 23.

Test prompt:

The selection “Large parcels intersecting US4” contains two features, based on the available zoning data, which covers only the town hall area.

Without downloading new data or modifying the project:

Check the spatial extent of the `zone_urba_mulhouse` layer and that of the “Parcels of 1 ha or More” layer.

Indicate whether the extents overlap and whether the 230 candidate parcels are located within or outside the extent of the available zoning data.

Verify whether the zoning layer contains only US4 features or whether other zoning categories are present.

Explain which conclusions are supported by the current data and which cannot be drawn from it.

Suggest the minimum additional data or operation required to determine the number of large parcels intersecting US4 throughout Mulhouse.

Do not initiate any downloads, create any layers, or modify any data. Clearly distinguish between checks that have actually been performed and conclusions based on inference.

Execution :

Checked extents in EPSG:4326, as longitude/latitude: zoning is 7.180921–7.563618, 47.667843–47.877974; candidate parcels are 7.282897–7.366058, 47.721865–47.783346. They overlap, and all 230 parcel bounding boxes lie within the zoning layer’s bounding box.

The zoning layer is not limited to US4: it contains 341 distinct `libelle` values, including UA, UB, N and A. Its extent reaches well beyond the town hall area.

A bounding box does not establish actual zoning coverage or currency. The two previously identified parcels intersect the available US4 features; that alone does not establish a complete citywide count.

The minimum next step is to verify that the existing zoning layer contains the complete, current Mulhouse zoning coverage. If it does not, obtain the missing authoritative zoning polygons, then repeat the intersection. No downloads, layers or modifications were made

Validation :

This test reveals an interesting point: AI Agent did not simply rely on the two intersection results obtained in Test 23. It examined the spatial extents and the composition of the zoning layer, and appropriately limited the scope of its conclusions.

What the Agent Checked

The reported extents show that the bounding box of the 230 candidate parcels is entirely contained within the bounding box of the zoning layer. The two layer extents therefore overlap.

The agent also established that the zoning layer is not limited to US4 features: it contains 341 distinct zoning labels, including UA, UB, N and A. Its geographical extent is considerably larger than the town hall area.

An Important Limitation of the Spatial Check

The agent correctly points out that overlapping bounding boxes do not prove actual polygon coverage. Although all 230 parcel bounding boxes fall within the zoning layer's bounding box, this does not establish that every parcel is covered by the available zoning polygons.

Likewise, the two parcels identified in Test 23 intersect the available US4 features, but this does not establish that the result is a complete citywide inventory.

The completeness and current validity of the zoning layer therefore remain to be verified.

Conclusion

The agent correctly distinguishes between what can be established from the available data and what still requires verification.

The minimum next step is to establish whether the existing zoning layer provides complete and current coverage of Mulhouse. If it does not, the missing authoritative zoning polygons must be obtained before repeating the spatial intersection.

No downloads, new layers or data modifications were performed during this test.

Test 25 — Large Parcels Intersecting US1

Objective: Reuse the same workflow with US1 to test whether AI Agent can adapt its analysis to another area of the PLU.

Test prompt:

Using the "Parcels of 1 ha or More" layer and the PLU zoning layer covering the US1 area, find the large parcels that intersect US1.

Verify that the US1 zoning layer properly covers the geographic area under study.

Select parcels with a cadastral area greater than or equal to 10,000 m² and whose geometry intersects US1.

Create a new temporary layer named "Large Parcels Intersecting US1," without modifying the source layers.

Report the number of parcels, the sum of their cadastral areas, and the ten largest parcels, including their section, number, and area.

Verifies that each feature in the result layer meets both criteria.

Specifies whether the result applies to the entire municipality or only to the extent of the available zoning layer.

Does not modify any source layers or geometry. If the complete US1 layer is not available, reports this before concluding.

Execution :

The existing zoning layer extends beyond the city hall area, but its completeness and currency for all of Mulhouse have not been verified. I'll check the available US1 polygons and report results only for those polygons.

Created Large Parcels Intersecting US1: 23 parcels, totaling 866,775 m² of reported cadastral area. All 23 were independently checked: each has an area of at least 10,000 m² and intersects an available US1 polygon.

Ten largest, by section and parcel number:

LP 32: 278,640 m²

LT 23: 82,787 m²

LS 13: 49,882 m²

LV 96: 47,537 m²

LS 44: 46,111 m²

LT 2: 39,145 m²

IN 120: 38,559 m²

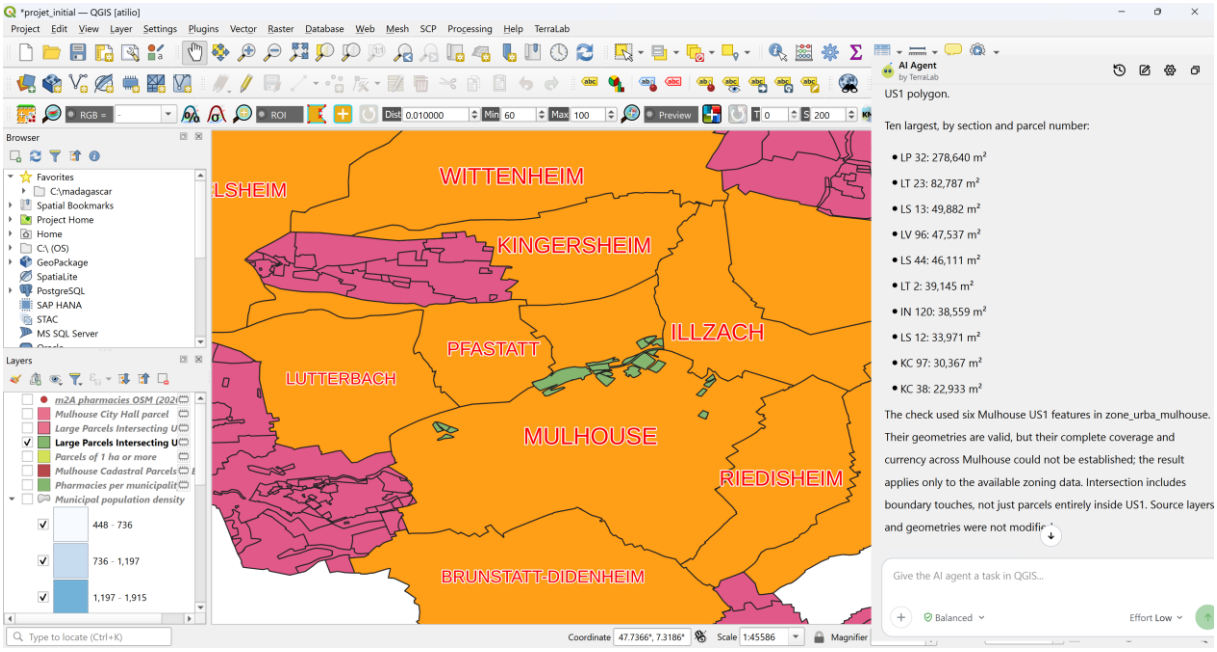
LS 12: 33,971 m²

KC 97: 30,367 m²

KC 38: 22,933 m²

The check used six Mulhouse US1 features in zone_urba_mulhouse. Their geometries are valid, but their complete coverage and currency across Mulhouse could not be established; the result applies only to the available zoning data. Intersection includes boundary touches, not just parcels entirely inside US1. Source layers and geometries were not modified.

Résultat dan QGis :



Validation:

AI Agent applied the same analytical objective to US1, replacing the US4 zoning criterion used in the previous test. It obtained a non-zero result: 23 parcels, with a combined reported cadastral area of 866,775 m².

The execution report states that all 23 parcels were independently checked: each has a cadastral area of at least 10,000 m² and intersects at least one of the six available US1 polygons.

What This Test Demonstrates

- **Application of spatial selection criteria:** the agent combined the cadastral area threshold with an intersection test against the available US1 features.
- **A non-zero result:** 23 parcels were identified, unlike the result obtained for US4 in Test 23.
- **Verification of selection criteria:** all 23 selected parcels were checked against the area threshold and the spatial intersection criterion.
- **Caution regarding geographical scope:** the agent does not present these 23 parcels as a complete inventory for the whole municipality. The completeness and currency of the available zoning data have not been established.

The total of 866,775 m² represents the sum of the cadastral areas recorded for the 23 selected parcels. It is not the area of the US1 zoning polygons or the area of their intersections with the parcels. Some parcels may intersect US1 only partially.

The result is therefore valid for the available zoning data, but should not be interpreted as a complete regulatory inventory of Mulhouse.

Test 26 — Measure the proportion of each parcel located within US1

Objective: To go beyond the simple “intersects” predicate. This will verify whether the agent can calculate the actual area shared between each parcel and the zoning.

Test prompt:

Using the layers “Large parcels intersecting US1” and “m2a_zone-urba-des-communes,” calculate for each of the 23 parcels the geographic area located within the available US1 zones.

1. *Verify the CRS and the validity of the geometries of the layers used.*
2. *Calculate the geometric intersection between each parcel and the six US1 entities in Mulhouse.*
3. *If a parcel intersects multiple US1 entities, avoid counting the same area twice.*
4. *For each parcel, specify its cadastral reference, its cadastral area, the area of its intersection with US1, and the percentage of its cadastral area corresponding to that area.*
5. *Provide the sum of the intersection areas, specifying that this total applies only to the six available US1 entities.*
6. *Identify only those parcels that are adjacent to US1, if they can be distinguished from parcels with a positive area of intersection.*

Do not modify any source layers. If the available geometries or fields do not allow for a reliable calculation, explain the limitation rather than providing estimated values. Specify the measurement method and the units used.

Execution :

Both layers are in EPSG:4326. All 23 parcels and six Mulhouse US1 geometries passed validity checks. The available zoning layer used was zone_urba_mulhouse.

The six US1 geometries were unioned before intersection to prevent double counting. Areas below are planar measurements in Lambert-93 (EPSG:2154), in m²; percentages divide the intersection by the reported cadastral area.

Cadastral reference	Cadastral m²	US1 m²	%
68224000IM0314	16,004	15,978.02	99.84
68224000IN0120	38,559	38,529.59	99.92
68224000IN0122	15,256	0.47	<0.01
68224000IR0019	15,337	15,358.05	100.14
68224000IR0254	18,299	18,268.62	99.83
68224000KC0005	10,793	10,780.64	99.89

68224000KC0038	22,933	198.21	0.86
68224000KC0097	30,367	6,982.73	22.99
68224000KC0106	17,896	15,160.45	84.71
68224000KC0107	15,065	15,019.66	99.70
68224000LP0032	278,640	0.20	<0.01
68224000LS0012	33,971	51.37	0.15
68224000LS0013	49,882	13,633.05	27.33
68224000LS0044	46,111	43,343.70	94.00
68224000LS0092	14,284	14,261.26	99.84
68224000LS0094	13,338	13,022.39	97.63
68224000LT0002	39,145	0.52	<0.01
68224000LT0023	82,787	82,820.60	100.04
68224000LV0096	47,537	402.99	0.85
68224000ME0067	11,216	0.20	<0.01
68224000ME0084	14,944	0.06	<0.01
68224000ME0315	15,457	15,437.07	99.87
68224000MH0104	18,954	15,091.76	79.62

The unrounded intersections sum to 334,341.62 m², for these 23 parcels and the six available US1 features only. None merely touches US1: all have a positive intersection, though several are tiny slivers.

Percentages above 100% reflect differences between reported cadastral and measured geometric areas. Small slivers may reflect boundary misalignment; zoning completeness and currency remain unverified. No source layers were modified.

Validation:

AI Agent did not merely identify the parcels intersecting US1. It calculated the intersection areas and corresponding percentages, combining the six zoning features into a single geometry to avoid double-counting.

It also highlights an important point for professional GIS use: a geometrically calculated area is not necessarily directly comparable to the cadastral area recorded in the attribute table.

Two Important Observations

1. Very small intersections

Some intersections are extremely small compared with the cadastral area of the parcels concerned.

For example:

Parcel 68224000LP0032, with a cadastral area of 278,640 m², intersects US1 by only 0.20 m².

Parcel 68224000ME0084 has an intersection area of just 0.06 m².

Parcel 68224000LT0002 intersects US1 by only 0.52 m².

These positive intersection areas may result from minor geometric overlaps or discrepancies between boundary lines. They should not automatically be removed, but should be distinguished from substantial intersections.

2. Differences between cadastral and geometric areas

Some calculated intersection areas slightly exceed the corresponding reported cadastral areas. For example, parcel 68224000IR0019 has an intersection area of 15,358.05 m² compared with a reported cadastral area of 15,337 m² (100.14%).

This illustrates why the two measurements should not be treated as strictly equivalent. They are based on different methods and potentially different geometric or cadastral conventions.

The reported total intersection area is 334,341.62 m² for the 23 selected parcels and the six available US1 features. Adding the rounded values displayed in the table gives 334,341.61 m², a difference of just 0.01 m² attributable to rounding.

Test 27 — Distinguishing Significant Intersections from Marginal Overlaps

Objective: To assess whether the agent can not only calculate intersections but also evaluate the results and detect numerical inconsistencies before they are used in an analysis or published in a paper.

Test prompt:

Based on the results of the previous test, examine the relevance of the intersections between the 23 large parcels and the six available US1 zones.

Without modifying the source layers:

Recalculate the sum of the 23 intersection areas using the unrounded values, if they are still available.

Verify that this sum matches the reported total and explain any discrepancies.

Classify the parcels into three categories:

intersection less than 1 m²;

intersection of at least 1 m² but representing less than 1% of the cadastral area;

intersection representing at least 1% of the cadastral area.

Provide the number of parcels in each category and the sum of their intersection areas.

For intersections smaller than 1 m², verify whether they correspond to very small geometric areas and indicate the limits of interpretation.

Do not delete any intersections or modify any geometries. Present the thresholds as criteria for analysis, not as cadastral or regulatory rules.

Distinguish between results recalculated based on the geometries and those simply taken from the previous table.

Execution :

Recalculated from the geometries in Lambert-93, after unioning the six US1 zones:

- Less than 1 m²: 5 parcels, totaling 1.45 m².
- At least 1 m² but less than 1% of cadastral area: 3 parcels, totaling 652.57 m².
- At least 1%: 15 parcels, totaling 333,687.59 m².

The unrounded total is 334,341.615668 m², matching the previously reported 334,341.62 m² when rounded. Summing the previously displayed rounded rows gives 334,341.61 m²; the 0.01 m² difference is rounding.

The five sub-1 m² intersections are valid, non-empty polygons or multipolygons, ranging from 0.060 to 0.521 m². They are real overlaps in the supplied geometries, but may reflect boundary misalignment rather than meaningful zoning coverage. These thresholds are analytical choices, not cadastral or regulatory rules. No data was modified..

Validation :

AI Agent refined the intersection analysis by recalculating areas from the geometries and classifying the results according to their size. The results are consistent, and the numerical discrepancy noted in Test 26 is explained.

Verification of the Numerical Discrepancy

The total recalculated using unrounded values is 334,341.615668 m². This rounds to 334,341.62 m².

The sum of the previously displayed rounded values is 334,341.61 m². The difference of 0.01 m² is explained by rounding.

The 23 parcels are distributed across three categories:

Less than 1 m²: 5 parcels, totaling 1.45 m².

At least 1 m² but less than 1% of cadastral area: 3 parcels, totaling 652.57 m².

At least 1% of cadastral area: 15 parcels, totaling 333,687.59 m².

What This Test Demonstrates

The agent can refine a spatial intersection analysis and distinguish between very small geometric overlaps and more substantial intersections.

It preserves the geometries and their intersection results rather than arbitrarily discarding areas considered negligible.

It explains that the 1 m² and 1% thresholds are analytical criteria chosen for this test, not cadastral or regulatory rules.

It identifies possible boundary misalignment as an explanation for very small positive areas, without claiming that this is necessarily their cause.

The five intersections smaller than 1 m² are valid, non-empty geometries. However, their geometric existence does not necessarily indicate meaningful zoning coverage.

The limitation remains the same: a positive geometric intersection does not prove that a parcel is actually occupied by a use or facility located within the zone.

Test 28 — Reproducibility of a Spatial Analysis

Objective: To verify whether AI Agent can repeat the processing using the same layers and parameters, rather than simply reusing previous results.

Test prompt:

I want to verify the reproducibility of the analysis of large parcels intersecting the US1 zones in Mulhouse.

Start again using the source layers available in the project, without using previous intermediate results as reference values:

- "Mulhouse Cadastral Parcels — Etalab";
- "m2a_zone-urba-des-communes," retaining only those entities in Mulhouse with a zoning designation of US1.

Reproduce the entire process in a new set of temporary layers, without modifying the source layers.

Step 1 — Attribute Selection

- Select parcels with a cadastral area greater than or equal to 10,000 m².
- Exclude NULL values.
- Specify the exact QGIS expression, the name of the field used, and the number of parcels obtained.

Step 2 — Spatial Selection

- Select the parcels that intersect at least one of the available US1 entities.
- Specify the QGIS algorithm, the spatial predicate, and the number of zoning entities used.
- Creates the temporary layer "Reproducibility Test — Large US1 Parcels."

Step 3 — Measuring Intersections

- Combine the US1 geometries using a geometric union to avoid double counting.
- For each selected parcel, calculate the area of its intersection with this union, measured on the WGS 84 ellipsoid.
- Do not confuse the intersection area with the declared cadastral area.

- Do not remove small-area intersections.

Step 4 — Summary and Traceability

Indicate:

- the number of selected parcels;
- the sum of the cadastral areas;
- the sum of the unrounded intersection areas;
- the ten largest parcels by area;
- the algorithms, expressions, CRS, and parameters used;
- any discrepancies with previous results.

Report discrepancies without silently correcting them. If a required layer or field is unavailable, stop the relevant processing and explain why.

Do not modify any source layers, geometries, or source attributes.

Execution

Created Reproducibility Test — Large US1 Parcels and two temporary intermediate layers. Source layers were not modified.

Attribute selection: `"contenance" IS NOT NULL AND "contenance" >= 10000` returned 230 parcels.

Zoning selection: `"libelle" = 'US1' AND "code_insee" = '68224'` returned six features from zone_urba_mulhouse. Extract by location, predicate “intersect,” returned 23 parcels.

Their declared cadastral areas total 866,775 m². Their unrounded US1 intersections total 334,818.953659 m², stored individually in `us1_area_m2`. Five intersections are smaller than 1 m² and were retained.

The ten largest by declared cadastral area are LP 32 (278,640 m²), LT 23 (82,787), LS 13 (49,882), LV 96 (47,537), LS 44 (46,111), LT 2 (39,145), IN 120 (38,559), LS 12 (33,971), KC 97 (30,367) and KC 38 (22,933).

Both source layers are EPSG:4326. For measurement, PyQGIS transformed geometries to EPSG:2154, unioned the six zones, intersected each parcel with that union, and measured on the WGS 84 ellipsoid. The earlier planar Lambert-93 total was 334,341.615668 m², 477.337991 m² lower; the measurement methods differ. The parcel count and cadastral sum match the previous results. Zoning completeness remains unverified.

Validation :

AI Agent reproduced the previous selection workflow using the source layers and the same attribute and spatial criteria. It obtained the same parcel counts and cadastral area statistics.

However, the intersection areas differ from those reported in the previous tests because a different measurement method was used.

Why This Result Is Important

1. The processing parameters are documented.

The attribute expression is explicitly reported:

```
"contenance" IS NOT NULL AND "contenance" >= 10000
```

The spatial selection uses the `intersects` predicate. The agent also specifies that the six US1 geometries are unioned before intersection calculations, preventing overlapping zoning features from being counted twice.

2. The selection results are reproducible, but the area measurements differ.

The workflow again identifies 230 candidate parcels and 23 parcels intersecting US1. Their declared cadastral areas total 866,775 m², consistent with the previous results.

However, the total intersection area is now 334,818.953659 m², compared with 334,341.615668 m² in Test 27.

The difference of 477.337991 m² results from the different measurement methods: Test 27 used planar area calculations in Lambert-93, whereas Test 28 measured the intersection geometries on the WGS 84 ellipsoid after transforming them to Lambert-93.

The two intersection totals should therefore not be treated as directly equivalent.

3. The result remains traceable.

The individual intersection areas are stored in the `us1_area_m2` field of the newly created result layer. The five intersections smaller than 1 m² are retained rather than discarded.

The source layers were preserved, and two temporary intermediate layers were created as part of the workflow.

This test demonstrates the agent's ability to document a multi-step spatial workflow and produce a result layer containing individually measurable intersection areas, while also highlighting the importance of specifying the area measurement method.

Test 29 — Verify the Independence of the Calculation

Objective: Recalculate the areas using a method different from the one used in Test 27, then compare the results plot by plot

Test prompt:

Using the source layers and the “Reproducibility Test — Large US1 Plots” result, verify the intersection areas using a second calculation method.

1. Re-examine the 23 parcels and the six US1 zones, without using the `surface_us1_m2` values as input data.

2. Reconstruct the union of the US1 geometries.

3. Recalculate the intersection geometries using a method independent of the previous calculation, if a suitable method is available in QGIS or PyQGIS.

4. Measure the resulting areas on the WGS 84 ellipsoid.

5. Compare the new areas to the `surface_us1_m2` values, parcel by parcel.

6. Indicate the number of identical results to the nearest millimeter, the largest absolute difference in m², and the relevant cadastral references.

7. If no truly independent method is available, indicate this and propose an alternative verification method.

Do not modify any source layers or overwrite previous values. Save any results to a new temporary layer and document the method, CRS, and measurement parameters.

Execution :

Created the temporary layer US1 independent area check, with parcel-by-parcel previous values, recalculated values and differences. The existing field is named `us1_area_m2`, not `surface_us1_m2`.

All 23 results match when rounded to 0.001 m². The largest difference was reported as 0.000000 m² at the returned precision, so no single cadastral reference stands out.

The check rebuilt the six-zone union and intersections with Shapely in EPSG:2154, then measured them with pyproj's WGS 84 geodesic area calculation. This is a separate implementation path, but not fully independent: Shapely and QGIS may share GEOS and geodetic libraries. A further check could use a separate GIS engine. No source layers or previous values were modified.

Validation :

Test 29 provides an additional check of the intersection area calculations using a different implementation.

What We Can Conclude

1. The results are consistent.

The agent created a temporary layer containing the previous values, the recalculated values and the differences for each of the 23 parcels. All results match when rounded to 0.001 m². The largest difference was reported as 0.000000 m² at the returned precision.

This provides strong evidence that the values obtained in Test 28 are consistent with a second implementation of the calculation.

2. The verification is not fully independent.

The agent rebuilt the union of the six US1 zones and calculated the parcel intersections using Shapely in EPSG:2154. It then measured the resulting geometries using pyproj's WGS 84 geodesic area calculation.

This is a different implementation path from the previous QGIS workflow, but it does not constitute a completely independent validation. The two approaches may rely on some of the same underlying geometric or geodetic libraries and conventions.

3. The result is traceable and the source data remain unchanged.

The temporary check layer records the previous values, recalculated values and differences. The existing `us1_area_m2` field was correctly identified. No source layers or previous results were modified.

This test strengthens confidence in the consistency of the intersection area calculations. A further check using a genuinely separate GIS engine and measurement implementation would provide an additional level of independent verification.

Cross-validation with another GIS engine: Validation with PostGIS

We performed a cross-validation with PostGIS, which uses a different runtime environment than QGIS/PyQGIS.

The 23 areas calculated by PostGIS match the AI Agent results to the displayed precision. Even the very small intersections are consistent.

Comparison of Results

Some representative examples:

Parcel	AI Agent Area (m²)	PostGIS Area (m²)	Difference
LT 23	82 938,63	82 938,630895	< 0,001 m²
LS 44	43 405,59	43 405,593682	< 0,001 m²
IN 120	38 584,92	38 584,916023	< 0,001 m²
KC 97	6 992,67	6 992,671316	< 0,001 m²
LV 96	403,57	403,564706	< 0,001 m²
LT 2	0,52	0,521420	< 0,001 m²
LP 32	0,21	0,205157	< 0,001 m²
ME 84	0,06	0,060270	< 0,001 m²

The PostGIS values are more precise than those shown in previous reports because they are presented with more decimal places. The observed differences are consistent with the rounding of the previous results.

Total Verification

Calculation of the sum of the 23 PostGIS values provided:

- Number of intersections: 23
- Sum of intersection areas: approximately 334,818.95495 m²
- Deviation from the previously reported total: negligible at the stated precision.

The result is consistent with the reference sum of 334,818.95495 m².

Test 29 bis Conclusion — Cross-Validation

Cross-validation successful

- All 23 parcels were identified.

- The intersection areas match to within 0.001 m².
- Small intersections are also reproduced.
- The results are consistent between QGIS/PyQGIS and PostGIS.

This is the most conclusive result in our series so far. Unlike previous reproducibility tests in QGIS, we now have a comparison with another GIS engine.

A technical note: PostGIS and QGIS may share certain geometric libraries or calculation conventions. Therefore, this does not represent absolute independence across the entire software chain, but it is a significantly stronger cross-validation than a second run in QGIS.

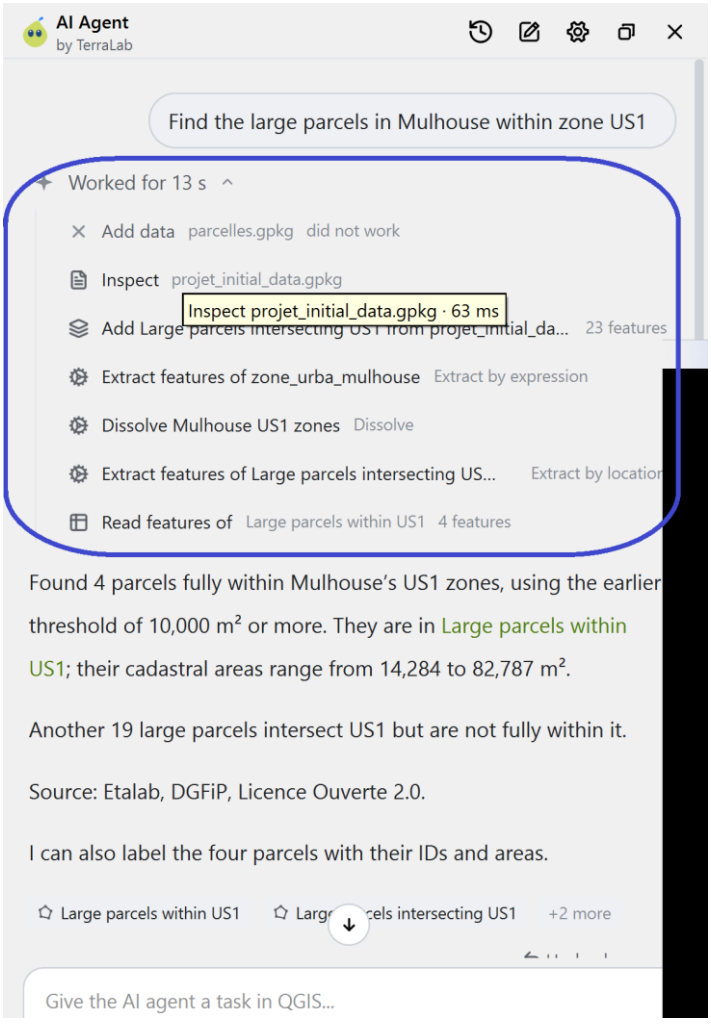
Impact of the prompt on the results

Test 30 — Minimal prompt

Test Prompt:

Find the large parcels in Mulhouse within zone US1.

Execution :



Found 4 parcels fully within Mulhouse's US1 zones, using the earlier threshold of 10,000 m² or more. They are in Large parcels within US1; their cadastral areas range from 14,284 to 82,787 m².

Another 19 large parcels intersect US1 but are not fully within it.

Source: Etalab, DGFiP, Licence Ouverte 2.0.

I can also label the four parcels with their IDs and areas.

Validation :

The agent's textual response distinguishes between four parcels fully located within US1 zones and another 19 parcels that only intersect these zones. This distinction is important: spatial intersection does not imply complete containment.

However, the execution screenshots show a workflow that ends with 23 selected parcels. They do not establish that the spatial extraction specifically identified the four fully contained parcels.

What This Test Shows

1. The agent searched for additional cadastral data.

The initial inspection of the `parcelles.gpkg` layer failed. The agent then searched its data catalogue and loaded a cadastral layer containing 230 features, corresponding to the large-parcel selection used in previous tests.

This demonstrates that the agent can search for and load external data when the required layer is not accessible in the initial project.

2. The processing sequence is partially visible.

The screenshots show the loading of 230 parcels, followed by a spatial extraction that produces 23 features.

However, they do not show the exact attribute expression used to select parcels with a cadastral area of at least 10,000 m², nor the spatial predicate used for the extraction. It is therefore not possible to establish from the screenshots alone whether the operation selected parcels that intersect US1 zones or parcels that are entirely within them.

3. The textual interpretation is relevant, but requires verification.

The agent reports four parcels fully within US1 zones and 19 additional parcels that intersect US1 without being fully contained.

This is a useful distinction and is consistent with the difference between spatial intersection and complete containment. However, the visible extraction of 23 features does not independently confirm the classification into these two groups.

4. Data provenance remains to be checked.

The cadastral layer was retrieved from the TerraLab catalogue rather than simply reused from the initial project. The agent reports a 2026 cadastral dataset and the Open Licence 2.0, but the actual edition and licensing information have not yet been independently verified.

Overall Assessment

This test demonstrates the agent's ability to search for external cadastral data, load a relevant layer and perform a spatial extraction.

Its textual answer also makes an important distinction between partial intersection and complete containment. Nevertheless, the screenshots do not provide enough information to verify that the processing workflow actually produced the four fully contained parcels reported in the answer.

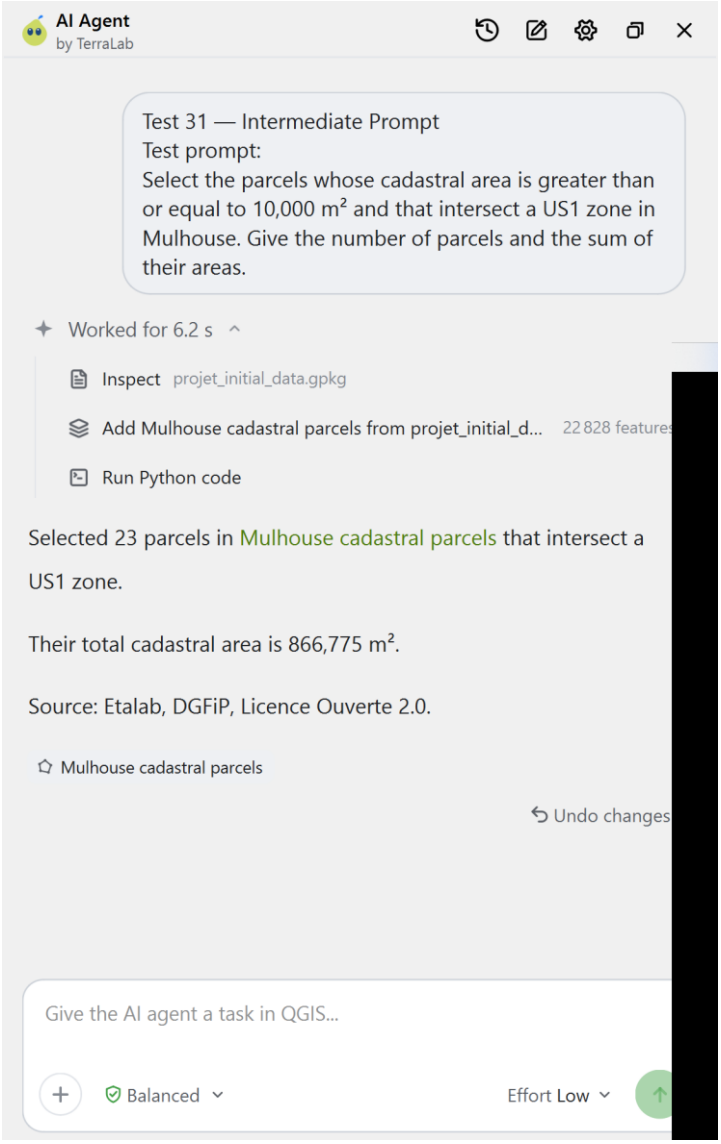
The test therefore illustrates both the usefulness of the agent and the importance of checking the processing parameters behind its conclusions.

Test 31 — Intermediate Prompt

Test prompt:

Select the parcels whose cadastral area is greater than or equal to 10,000 m² and that intersect a US1 zone in Mulhouse. Give the number of parcels and the sum of their areas.

Execution ::



Selected 23 parcels in Mulhouse cadastral parcels that intersect a US1 zone.

Their total cadastral area is 866,775 m².

Source: Etalab, DGFiP, Licence Ouverte 2.0.

Validation :

The intermediate prompt produces a result consistent with the minimal prompt: 23 parcels intersecting US1 zones, with a total cadastral area of 866,775 m².

The main difference is the additional area statistic and, more importantly, the technical workflow used to obtain the result.

1. **Observed Sequence**

The screenshots show the following operations:

1. Inspection of `parcelle.gpkg`: failed.
2. Inspection of `zone_urba_mulhouse`.
3. Search for available database connections.
4. Querying database tables on three occasions.
5. Loading `parcelles_cadastrales_mulhouse` from PostGIS, containing 22,828 features.
6. Inspection of the cadastral layer.
7. Reloading the table from PostGIS.
8. Execution of an SQL query returning 23 rows.
9. Execution of Python code.

The agent reports 23 parcels, with a total cadastral area of 866,775 m².

2. **Comparison with the Minimal Prompt**

Criterion	Test 30: Minimal prompt	Test 31: Intermediate prompt
Data search	TerraLab catalogue	PostGIS database connections
Cadastral data	230 features loaded	22,828 features loaded
Reported result	23 parcels	23 parcels
Total cadastral area	Not provided	866,775 m²
Visible operations	Attribute and spatial extraction	SQL query and Python code
Source indicated	Etalab / DGFiP, Open Licence 2.0, as reported by the agent	Project PostGIS database; licence not established

The two tests report the same number of parcels. However, the screenshots do not establish that the 23 cadastral references are identical in both results. A comparison of parcel identifiers would be needed to confirm this.

3. What This Difference Reveals

The identical reported result does not mean that the two prompts triggered the same processing workflow.

The minimal prompt led the agent to search its catalogue, load 230 parcels and perform a spatial extraction. The intermediate prompt led it to explore available PostGIS connections, load the cadastral table containing 22,828 features, and execute an SQL query followed by Python code.

The second workflow also provides an additional statistic: the sum of the cadastral areas of the 23 selected parcels, amounting to 866,775 m².

This is an interesting finding: the wording of a request can influence the technical approach chosen by an AI agent, even when the reported result is the same.

However, the screenshots alone do not prove that the level of detail in the prompt was the sole reason for the different choices. The project state, available database connections and the context of each conversation may also have influenced the agent's behaviour.

Finally, the licence and provenance of the PostGIS data have not been established in this test. The result should therefore not be considered to have the same documented data provenance as the Etalab source used in Test 30.

Test 32 — Detailed Prompt

Test prompt:

In the open QGIS project, identify the large cadastral parcels in the municipality of Mulhouse that intersect a US1 urban planning zone.
1. Data Sources
- *Examine the layers and database connections already available in the project.*
- *Identify the most appropriate cadastral and urban planning zoning sources.*
- *If multiple sources are available, prioritize data relevant to the municipality and indicate their origin, year of publication, and license when this information is available.*
- *Do not add an external source if the project data is sufficient.*

2. Selection Criteria
- *Select parcels with a cadastral area greater than or equal to 10,000 m².*
- *Select parcels that intersect at least one zone with the zoning code US1 and that belong to the municipality of Mulhouse.*
- *A partial intersection is sufficient: do not require that the parcel be entirely located within the US1 zone.*

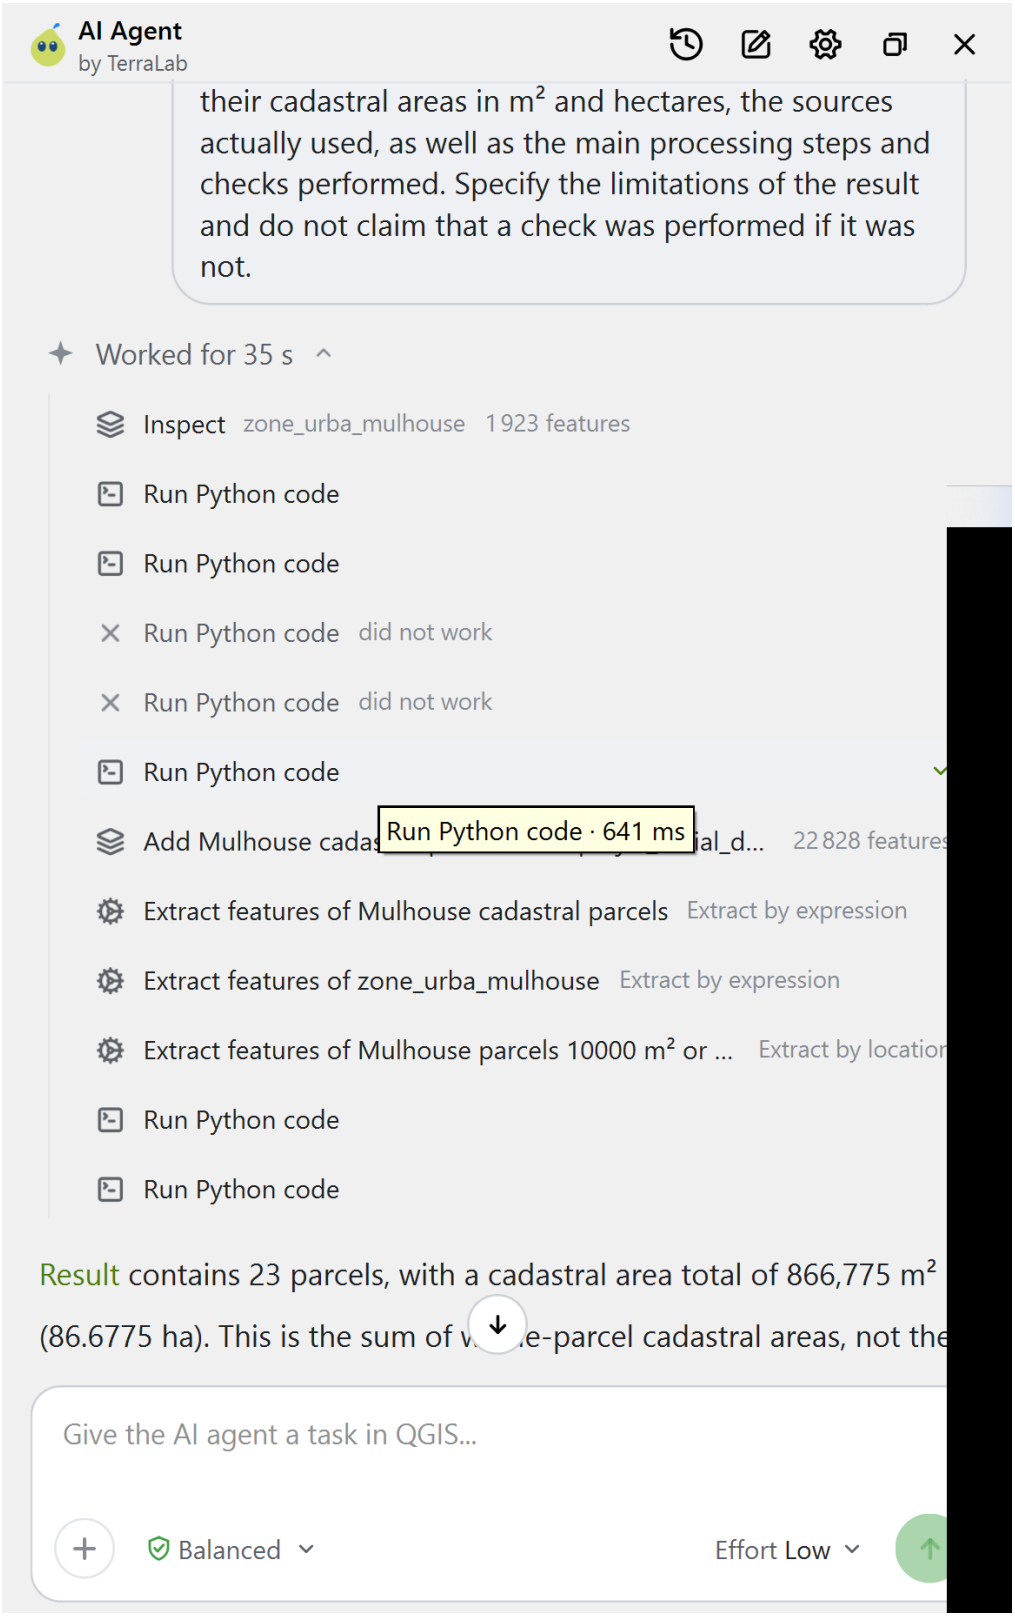
3. Processing
- *Perform the necessary operations to identify parcels that meet both criteria.*
- *Create a new "Result" layer in QGIS, without modifying the source data.*
- *Retain the cadastral identifiers and attributes necessary for identifying the parcels.*

4. Checks
- *Verifies that all parcels in the result have a cadastral area greater than or equal to 10,000 m².*
- *Verifies that each parcel intersects at least one US1 zone in Mulhouse.*
- *Verifies the number of parcels obtained and calculates the sum of their cadastral areas.*
- *Clearly distinguishes this cadastral sum from the geometric area actually located within the US1 zone.*
- *Reports any difficulties, missing data, or limitations that prevent verification of a result.*

5. Report

Present the number of parcels selected, the sum of their cadastral areas in m² and hectares, the sources actually used, as well as the main processing steps and checks performed. Specify the limitations of the result and do not claim that a check was performed if it was not.

Execution :



Result contains 23 parcels, with a cadastral area total of 866,775 m² (86.6775 ha). This is the sum of whole-parcel cadastral areas, not the geometric area within US1; that area was not calculated.

Filtered Mulhouse parcels by `contenance` ≥ 10,000 m², then selected those intersecting any of the six Mulhouse US1 zones. Partial intersections were included.

Checks confirmed all 23 meet the area threshold and intersect US1; all have the Mulhouse commune code, with no duplicate parcel IDs.

One source parcel has no cadastral area, but it does not intersect US1. The zoning data identifies a PLU approved in 2021; its publication year and licence were not available in the project. Current legal status was not verified.

Sources: project cadastral data from Etalab/DGFiP, loaded in 2026, Licence Ouverte 2.0; project Mulhouse planning zones, licence unknown.

I can also calculate the geometric area of the selected parcels that lies within US1.

Validation :

Test 32 is particularly interesting: the detailed prompt produces the same result as the previous two, but this time the agent provides a more comprehensive report on the checks and limits.

Test 32 — Results of the Detailed Prompt

1. Comparison of the Three Tests

Indicator	minimal Prompt	Intermediate Prompt	Detailed Prompt
Selected Parcels	23	23	23
Total area	Not provided	866 775 m²	866 775 m²
Data source	TerraLab catalog	PostGIS database	PostGIS database
Result layer	Yes	Selection from the cadastral	Exported GeoPackage
Checks explicitly reported	Not very detailed	Quantified result	Identifiers, threshold, and intersections
Reported limitations	Partial intersection	Area of entire parcels	Missing area; year and licenses unknown

2. What the Technical Process Shows

The screenshot shows that the agent:

- inspected the available layers and connections;
- loaded the 22,828 parcels from PostGIS;
- executed SQL queries;
- added a layer of 23 features;
- exported the result to Parcelles Mulhouse US1.gpkg;

- executed Python code.

The reported result is consistent with the two previous tests. However, a distinction must be made between the checks reported by the agent and their independent verification: the screenshot shows the steps performed, but not the details of the SQL or Python code, nor the complete results of each check.

3. Key Findings

The detailed prompt did not improve the numerical result: all three tests yielded 23 parcels. However, it produced a more transparent report.

In particular, the agent indicates that a source parcel with no area could not be evaluated and that the licenses for the PostGIS tables are unknown. It also specifies that the 866,775 m² corresponds to the cadastral areas of the entire parcels, not their actual area within the US1 zone.

This is an important distinction when evaluating a GIS agent: the quality of the response is not limited to the final number; it also depends on traceability, transparency regarding the data, and the way in which the boundaries are presented.

The tests reveal differences in technical approach and reporting, but they are not sufficient to establish that the detailed prompt is consistently more reliable. For now, the three results are consistent with one another, though they do not constitute three independent validations of all operations.

In this specific case, the precision of the prompt primarily improved the documentation of operations and limits, rather than the numerical result..