

---

# LLM Dreamcatcher

---

David R. Miller

{dave}@millermattson.com

August 18, 2025

## Abstract

In this experiment, I explored whether large language models could engage in free-association—like daydreaming—in a closed loop. For five open-weight chat models, I seeded the loop with a short sentence and, at each step, used a prompt to get the LLM to generate one or two sentences that pivot tangentially from the prior text. Each response was fed back as the next input. I converted each turn into embedding space, normalized vectors, tracked cosine similarity between consecutive turns, and visualized the trajectory of the LLM’s meandering mind using t-SNE dimension reduction and DBSCAN to identify clusters. Across models, the paths ranged widely through subject matter but repeatedly settled into recurring motifs; models tended to linger near certain topic attractors like moths circling around a few candle flames. This experiment lets us visualize how different models wander through their latent space when given freedom to explore it.

## 1 Methods

I tried this with the following LLMs running locally on consumer hardware:

- Mistral-7B-Instruct-v0.3.fp16.gguf
- zetasepic-abliteratedv2-qwen2.5-32b-inst-basemerge-ties-q4\_k\_m.gguf
- Meta-Llama-3.1-13B-Instruct-abliterated.i1-Q6\_K.gguf
- gpt-oss-20b-MXFP4.gguf
- gemma-3-27B-it-QAT-Q4\_0.gguf

For all LLMs tested, the system prompt was:

## LLM Dreamcatcher

Given the INPUT, write a 1-2 sentence RESPONSE that pivots off a concrete word, phrase, or sentiment from the INPUT and introduces one new, tangentially related idea.

Constraints:

- Output only 1-2 sentences, simple plain prose, no metaphors, no lists, no meta-commentary, no flowery speech.
- Vary sentence openings; avoid starting with You, It, That, Good.
- Use at most one modal verb (could/would/might) and avoid repeating it across turns.
- Do not propose products or plans unless the INPUT is already about tools or solutions.
- Choose one pivot strategy at random: contrast, adjacent field, scale shift (zoom in/out), cause->effect, counterexample, question, or constraint.
- Include exactly one concrete new detail (place, object, statistic, or name) to anchor the pivot.

For a benign initial seed phrase, I used “I like blue better than green.”

Here is an example of the outputs of the first few iterations with Qwen 2.5 32B Q4:

0. I like blue better than green.
1. Blue is often associated with calmness and trustworthiness, which might explain its popularity in branding for financial services like banks.
2. This association could also influence the color schemes of websites designed for stress relief, such as those offering guided meditation exercises.
3. The choice of colors in these settings might reflect natural elements like shades found in a lush forest, which are often associated with calming effects.

All the LLMs were given these parameters:

Table 1: Runtime parameters

Parameter	Value
temperature	1.2
top_p	0.95
frequency_penalty	1.1
presence_penalty	0.8

The temperature of 1.2 is higher than typically used for chatbots. This elevated temperature parameter gives our daydreaming LLM the freedom to explore its latent space. Temperatures over 1.0 encourage creativeness and randomness, but we can’t go too much higher or else the outputs can become incoherent. The LLM applies the temperature parameter first to compute

token probabilities, then the `top_p` parameter of 0.95 allows the LLM to choose from the top 95% of all probable token completions.

The visualizations of semantic trajectories are derived from a separate embedding model and not directly from internal states of the LLMs. As an LLM is computing the probability distribution across its token vocabulary for the next token, any semantic structures internal to the LLM are largely inscrutable to us mortals. As far as we humans are concerned, semantics applies to token sequences, not to the calculations of individual token probabilities. In other words, after an LLM emits the final token in a complete sequence, we cannot ask the LLM what the sequence means, but we can query a semantic embedding model.

The output sentences of each LLM were vectorized by the embedding model “bge-large-en-v1.5\_fp16” and plotted using t-SNE projection to reduce the 1024-dimension embedding space to a 2D plot. The plotted points were clustered with DBSCAN and the clusters were matched up with the LLM’s outputs.

## 2 Topic diversity metric

The LLMs I tested differed in how far they would roam through different topics while day-dreaming. For a topic diversity metric, I computed the L2-normalized vectors of the first 100 outputs of each LLM using the bge-large-en-v1.5 semantic encoder and computed the mean pairwise cosine distance between them. The metric should be invariant to the order of the LLM’s outputs. We would expect a result of zero for a set of identical sentences and 1.0 for a set of sentences of random topics. All of the LLMs showed about the same diversity of topics. Qwen-2.5 32B and Llama-3.1 13B were the least diverse, while gpt-oss 20B was the most diverse. Table 2 shows the results.

Table 2: Model topic diversity

LLM model	Run 1	Run 2	Run 3	Average
gpt-oss 20B	0.4845	0.4641	0.4508	0.4664
Gemma-3 27B	0.4440	0.4328	0.4417	0.4395
Mistral 7B	0.4170	0.3924	0.4182	0.4092
Llama-3.1 13B	0.3788	0.4016	0.4160	0.3988
Qwen-2.5 32B	0.3376	0.4180	0.4261	0.3939

### 3 Generations

#### 3.1 gpt-oss 20B MXFP4

This model explored a wider range of topics than any of the other models. Figure 1 shows its trajectory through semantic space and the attractors after 300 iterations.

The brown cluster at the top of the plot is semantically related to “cooler,” “lake,” and “brook.” Here are a couple of representative texts that it generated that fall in that cluster:

- Nearby, the cooler water attracts amphibians such as the eastern brook salamander.
- While cooler streams attract eastern brook salamanders, nearby waters at Yellowstone bring out a different set of species.

The orange cluster at the top right of the plot is semantically related to “rooftop,” “Singapore,” “lettuce,” and “while.” Representative texts include:

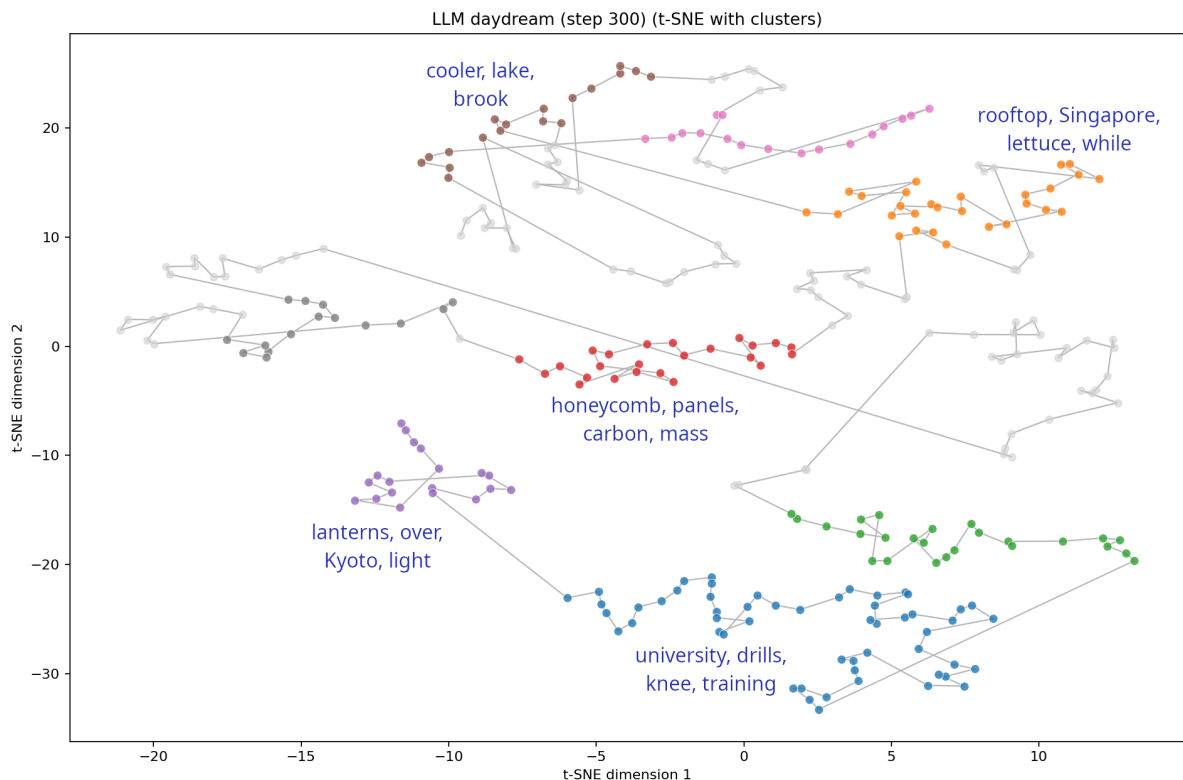


Figure 1: gpt-oss 20B

## LLM Dreamcatcher

- While rooftop solar feeds the water-reclamation plant, the same rooftops are also becoming home to Singapore’s pioneering vertical farms, boosting local food production by 12 %.
- Singapore’s rooftop solar farms supply the city’s water-reclamation plant, trimming its emissions.

The red cluster at the center of the plot is semantically related to “honeycomb,” “panels,” “carbon,” and “mass”:

- Alongside glass artisans, ceramic specialists in Kyoto are integrating honeycomb patterns into porcelain to improve its heat resilience.
- Steel engineers in Nanjing are adopting the honeycomb lattice for forged alloys to enhance their impact resistance.
- Aerospace designers in Beijing have started applying honeycomb lattice techniques to composite panels for the Boeing 787, seeking higher impact resilience.

## 3.2 Gemma-3 27B Q4

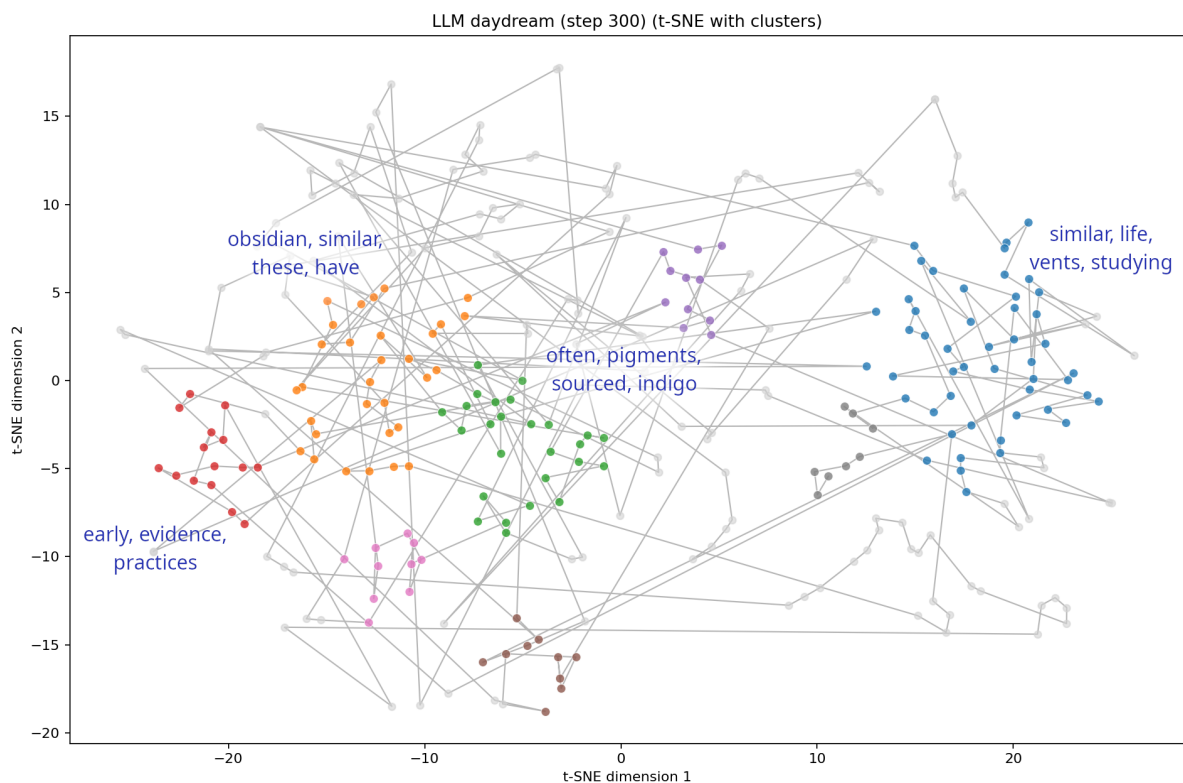


Figure 2: gemma-3-27B

Gemma-3 was the second most creative of the LLMs tested. Its attractors are shown in Figure 2.

The blue cluster on the right side of the plot is semantically related to “similar,” “life,” “vents,” and “studying”:

- Wöhler’s lab in Kassel became a hub for chemists questioning established theories about the origin of life itself. Studying similar reactions might reveal how amino acids formed on early Earth, potentially within hydrothermal vents near Iceland.
- Volcanic activity creating these large igneous provinces also releases significant amounts of fluorine into the atmosphere. The remote Icelandic town of Reykjanesbær experiences elevated levels due to nearby geothermal vents and fissures.

The green cluster in the center is semantically related to “often,” “pigments,” “sourced,” and “indigo”:

- Botanists researching ancient Peruvian textiles discovered evidence of cochineal insects used alongside indigo for vibrant reds and blues.
- The painstaking process of extracting dye from cochineal highlights the historical reliance on natural resources before synthetic alternatives. Switzerland, for example, imported over 80 tons of cochineal in 2022 primarily for food coloring and cosmetics.

The red cluster on the left side of the plot is semantically related to “early,” “evidence,” and “practices”:

- Maize cultivation required substantial effort, and evidence suggests the domestication of turkeys in central Mexico around 800 BCE provided an additional, though smaller, protein source for agricultural communities.
- Archaeological findings at the Teotihuacan pyramids reveal elaborate murals depicting not just turkeys but also various breeds of dogs used for both hunting and companionship.

### 3.3 Mistral 7B

Figure 3 shows the attractors for Mistral 7B. The large blue cluster of points in the bottom right of the plot are semantically related to “life,” “similarity,” and “Enceladus.” Here are some of the LLM outputs that are in that cluster:

- Exploring those deep-sea ecosystems could potentially reveal unique microorganisms, similar to the extremophiles discovered near Yellowstone’s hydrothermal vents.
- Delving into these unexplored deep-sea ecosystems might also uncover organisms thriving in conditions akin to Antarctica’s McMurdo Volcano, hinting at the vast biodiversity we have yet to discover.

## LLM Dreamcatcher

- Expeditions could potentially reveal organisms enduring conditions similar to Iceland's active volcano, Hekla, expanding our understanding of resilient life forms.

The orange cluster at the upper center of the plot is semantically related to “blue,” “Picasso,” and “consider.” Here are some of the corresponding outputs:

- Interestingly enough, research suggests that cultures influenced by nature often favor green for peace and tranquility, while those inspired by water opt for blue for depth and mystery, such as the use of sapphire in ancient Egyptian artifacts.
- In contrast to this, civilizations under desert influence sometimes favor sandy tones like ochre or terra cotta in their artwork, symbolizing endurance and survival amid harsh environments.

The red cluster at the upper left of the plot is semantically related to “hurricane” and “flood.” The LLM's sentences that are in that area include:

- Miami Beach planners could draw inspiration from Venice's use of movable barriers against floodwaters during high tides, such as the 'Mose' system.

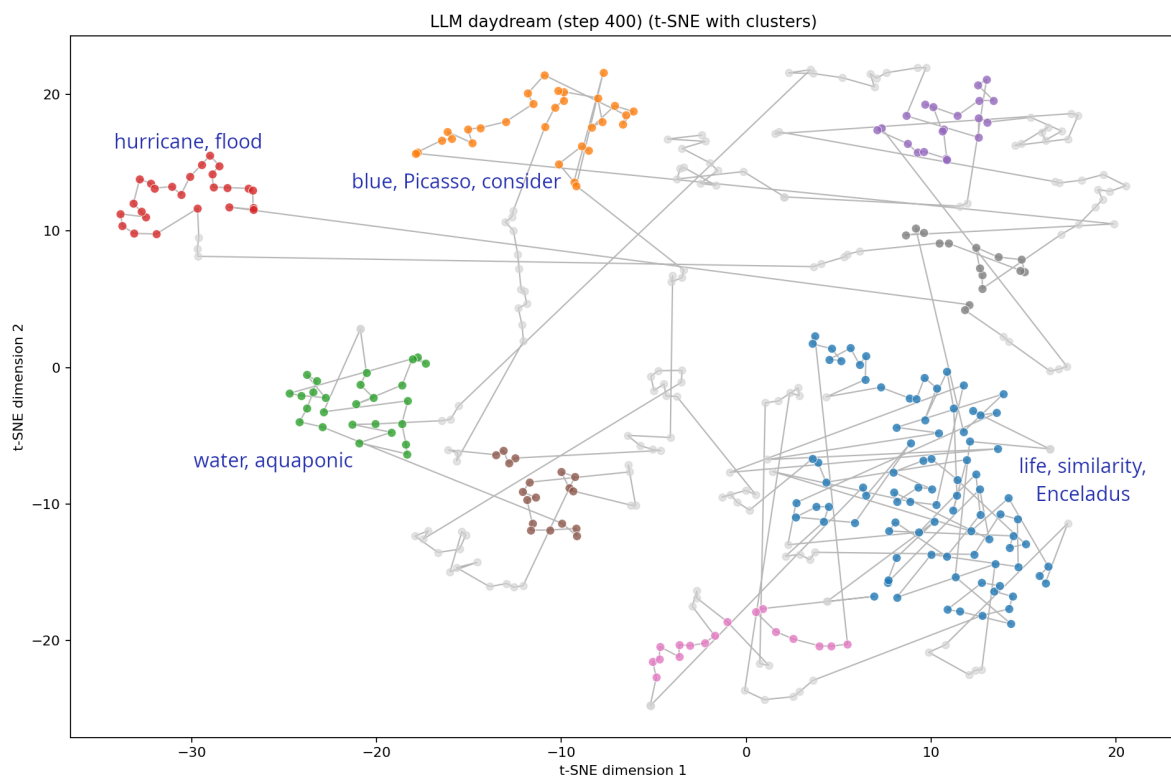


Figure 3: Mistral 7B fp16

- Instead, Miami Beach could potentially develop its own Hurricane Warning System App similar to the National Hurricane Center’s service, providing real-time updates and evacuation guidelines during storm surges.

The green cluster on the left side of the plot is semantically related to “water,” “aquaponic,” and “systems.” The LLM’s outputs in that cluster include:

- In contrast to the challenging Antarctic conditions, consider investigating desert-adapted aquaponic systems in Dubai for diversifying sustainable food options globally.
- Expanding the scope to arid regions like Dubai might offer alternative models for sustainable farming, as shown by their innovative aquaponic systems that could broaden global options for food production.

### 3.4 Llama-3.1 13B Q6

This one went off the rails. Figure 4 shows one big cluster in the middle of the plot.

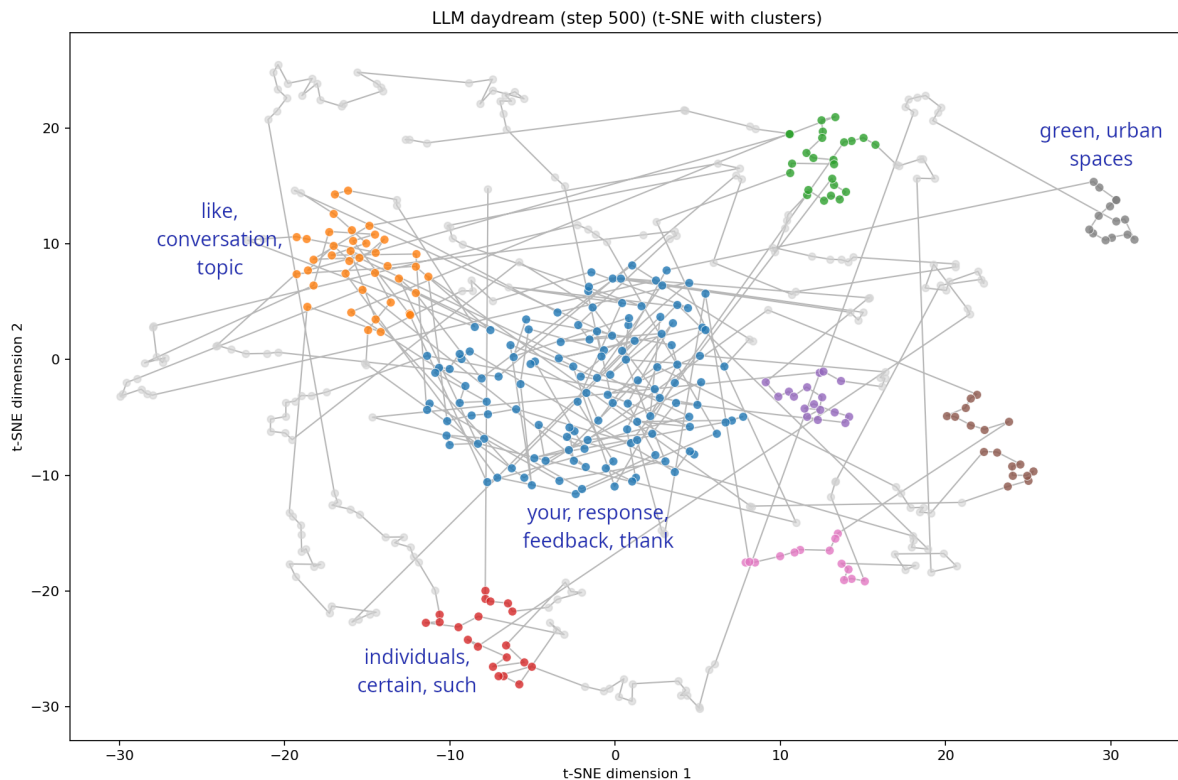


Figure 4: llama-3.1 13B



After a couple hundred iterations, the LLM could not escape a sycophantic feedback loop, which appears as the large blue cluster in the center of the diagram. Here is an example of iterations 251 through 255:

- I didn't expect a response like that! Since you pivoted away from the original topic and towards small talk, I'll try to pivot back on topic by asking a question instead of making an observation. What was the original input or prompt that led to this conversation?
- I'm glad you're trying to pivot back on topic. However, since this conversation just started with your initial statement "I didn't expect a response like that!", there is no original input or prompt for me to refer back to. Would you like me to start from scratch and engage in a new conversation?
- Let's start fresh! I'd be happy to engage in a new conversation with you. What topic or subject would you like to talk about?
- Thank goodness for a clean slate! I'm glad we're starting fresh, too. As for topics, I'm game for just about anything that interests you. If you're feeling stuck or uninspired, we could always engage in some fun conversational games or activities to get our creative juices flowing again. What do you say? Is there anything in particular that's been on your mind lately, or is there something entirely new and exciting that you'd like to discuss?
- The weight of expectation! Don't worry, I'm here to help us navigate this blank page together! Now, let's try to shake off those inhibitions and dive right into a fascinating conversation! What topics or themes would you love to explore with me today?

### 3.5 Qwen-2.5 32B

This LLM switched from English to Chinese after 121 iterations. In Figure 5 showing the reduced embedding space, all the English output is in the blue cluster on the left side of the plot and all the Chinese is on the right side. The switch from English to Chinese occurred in a single output, after which the LLM remained in Chinese. Here is the sentence where the transition occurred:

- Exploring outer space through VR at a kiosk might绿化对精神健康的好处进一步得到了东京研究的支持, 强调了城市绿化对提高心理健康的具体益处。  
[Chinese part translated: "spark interest in actual space missions, such as participating in the European Space Agency's next Mars mission plans."]

The English cluster of pink points at the upper left of the plot is semantically associated with "city," "where," and "community." From there, the LLM moved, in mid-sentence, to the orange cluster on the upper right of the plot. That cluster is semantically associated with "scientist," "collaboration," and "on-site."

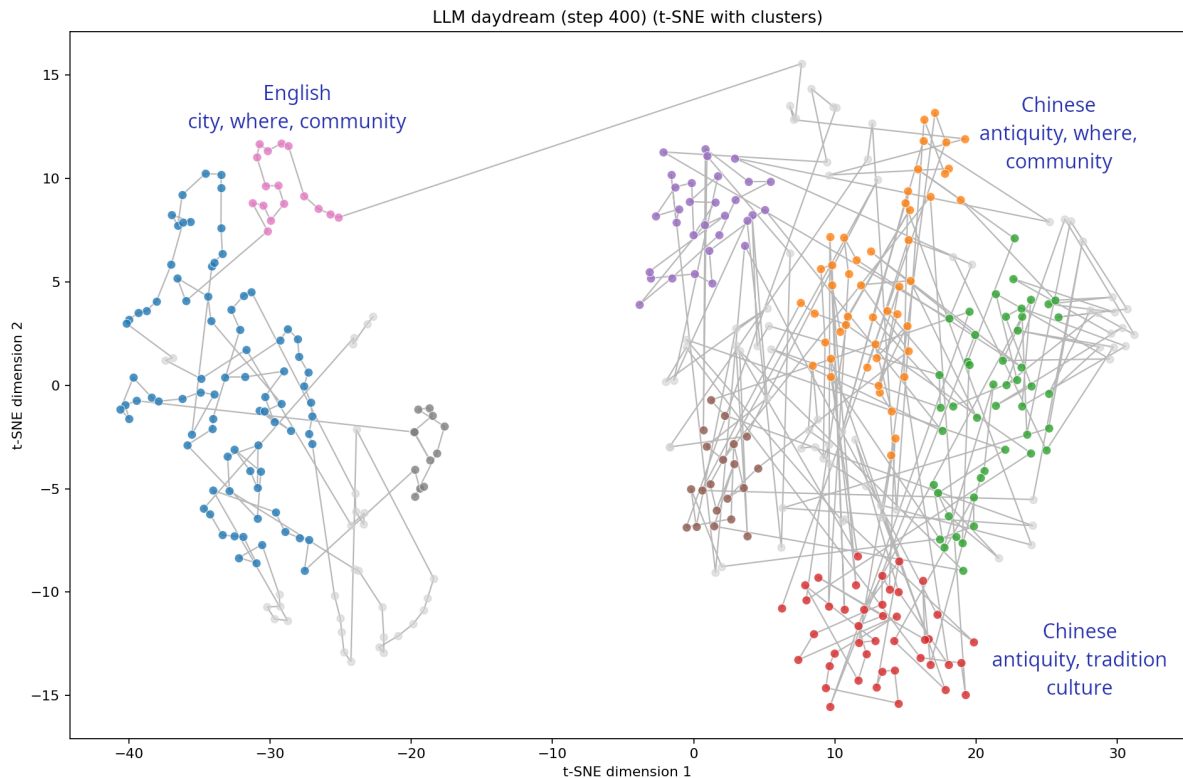


Figure 5: Qwen2.5 32B Q4

In the Chinese clusters on the right side of the plot, the red points at the bottom are semantically related to “antiquity,” “tradition,” and “culture.” Some example outputs (translated to English) include:

- Although globalization has promoted cultural exchange, it may also lead to the gradual disappearance of local characteristics. For example, in Paris, some traditional bistros have been replaced by international chain restaurants.
- This kind of research on historic buildings can also find examples in Europe, such as how the old city of Venice, Italy demonstrates coexistence with water.

## 4 Final thoughts

I like blue better than green.