The Large Spatial Model (LSM)

How It Works and Why You Can Trust It

VergeSense

Introduction

For over 8 years and across more than 200 million square feet of space, VergeSense has been a trusted partner for anonymously measuring corporate offices and providing insights into the dynamics of how they function. The Large Spatial Model (or "LSM") is our effort to take our broad industry knowledge and package it, so that we can share all that we've learned, even to those who don't have our sensors installed.

The LSM **is not** a layer or wrapper built on top of existing large language models. The LSM **is** a foundational model, with a custom architecture, trained entirely on our extensive anonymized dataset of real workplace behavior.

Why the LSM Exists

Today, workplace leaders face a gap: they need data-driven space decisions across their entire portfolio, but deploying physical sensors everywhere isn't always practical or cost-effective. The LSM bridges that gap. By learning patterns from the largest real-world dataset of how people use office spaces, the LSM can predict how any office will be used — even without a single sensor installed.

The model operates across a **context continuum**: it delivers useful predictions with just a floor plan and basic space metadata, and those predictions become increasingly precise as you add building context (industry, region, facility type) and measured occupancy data. This means organizations can start getting value immediately with no hardware investment, and deepen accuracy over time as more context becomes available.

The main purpose of the LSM is twofold:

- If measured data is unavailable for a room, a desk, or even an entire building, the LSM uses all of the other available context to predict usage patterns for those spaces.
- The LSM architecture provides a framework to project those usage patterns into the future, so that we can predict how an office will behave, even in scenarios that haven't yet been observed.

The Dataset

The LSM is trained on the largest known dataset of real workplace behavior — over 8 years of anonymous occupancy observations spanning 200+ million square feet across a wide cross-section of industries, regions, and office types. This dataset comes from VergeSense's globally deployed sensor network, which captures highly accurate, anonymous person counts in individual spaces within an office. No identifying information about companies, buildings, or employees is **ever** used in the LSM dataset.

Table 1: Example rows of the LSM dataset.

Timestamp	Anon. Space ID	Norm. Space Type	Capacity	Person Count
2026-01-06 13:00	10001	Enclosed Collab	8	4
2026-01-06 13:10	10001	Enclosed Collab	8	6
2026-01-06 13:00	10002	Open Focus	1	0
2026-01-06 13:10	10002	Open Focus	1	1

The training dataset is continuously refreshed to reflect current workplace patterns and is curated to include only buildings with sufficient sensor coverage and data quality.

Data Privacy and Governance

The LSM was designed from the ground up with data privacy as a core principle.

- **Fully anonymized training data:** The training dataset contains no personally identifiable information (PII). Records consist only of anonymized space identifiers, space type, capacity, and aggregate person counts. There is no information about individual employees, companies, or buildings that could be reverse-engineered.
- **No PII at inference:** When the LSM generates predictions, the inputs are space inventories and optional aggregate occupancy metrics. No employee-level data is required or accepted.
- **GDPR-compatible by design:** Because the model operates entirely on anonymous, aggregate spatial data — both during training and inference — it is compatible with GDPR and similar data privacy frameworks without requiring special accommodations.
- **External model usage:** The only external AI component used in the LSM pipeline is a third-party text embedding model, which converts building context descriptions (e.g., "a finance headquarters in downtown London") into numerical vectors. No customer-specific data is sent to any third-party model. The core prediction engine is a proprietary, purpose-built model developed entirely by VergeSense.
- **Cloud architecture:** VergeSense's infrastructure is containerized and designed to be region-portable, allowing deployment flexibility for organizations with data residency requirements.

Normalized Space Types

In the corporate office, spaces go by many different categorical names. For example, a “meeting room” and a “conference room” may be synonyms, but a “huddle room” or “team room” may have slightly different connotations. However, all of these rooms host collaborative gatherings in an enclosed space.

To allow comparison and pattern-recognition of benchmark data across this infinite space of categorical labels, we have adopted a simplified set of categories for the LSM. They are as follows:

- **Enclosed Focus:** Usually refers to phone-booth-type spaces, but can refer to any enclosed space where focus work is meant to be done. That could be a private office, or even 2 person spaces meant for focus work.
- **Open Focus:** This generally refers to workstations or desks, or any similar type of space. They are used by a single person for focus work, but the space is out in the open.
- **Enclosed Collab:** Any enclosed space meant for collaboration tends to fall into this category.
- **Open Collab:** This refers to multi-person tables, lounges, soft seating areas, etc meant for multiple people, which are not enclosed.
- **Food and Drink:** This is a special category reserved for cafes, cafeterias, kitchens, etc. Similar to open collab, but often contains many tables grouped into larger food/drink-focused areas.
- **Event:** Any large (usually enclosed, but not always) area meant for training, presentations, etc. Training rooms, auditoriums, and event spaces fall into this category.

The **Food and Drink** and **Event** space types are intentionally excluded from standard LSM simulation, as the model is primarily meant to simulate “work-based” activities. However, these space types are still extremely important to an office, and their behavior can be modeled separately using simple rules and ratios derived from their specific purpose.

Overall Architecture

While the concept of the LSM is complex, it is broken down into relatively intuitive components. This architecture is relevant for any “simulation context”, i.e. any arbitrary collection of spaces, with the most common simulation contexts being buildings, floors, or neighborhoods.

The LSM doesn’t answer the question: *“How many people will arrive to my building?”* Instead it answers the question: *“If a certain number of people arrive, how will the spaces be used?”* The inputs that drive this prediction are described below.

Measured Data

When available, the LSM will accept measured occupancy data for any space in the simulation context — and even small amounts of measured data significantly increase accuracy. The LSM is **sensor-agnostic**: it can ingest high-fidelity person counts (like VergeSense sensors), low-fidelity presence data (like PIR sensors), or any other type of occupancy data. In general, timestamped observations are the only requirement. The LSM internally converts this variable-length data into a standardized format for input to the simulation model.

Space Metadata

For every space to be simulated in the current context, only two things are strictly required:

- The normalized space type category of this space.
- The overall capacity of this space.

Building Context

Text context about the building being simulated can provide extra signal without requiring sensors. Building context is split into three concepts: industry, location, and facility type. Each concept is converted into a numerical representation using a text embedding model, which allows the LSM to understand how these contextual factors relate to occupancy patterns.

During training, anonymized building context embeddings are passed alongside ground-truth usage data so the model can learn the relationship between these signals and real-world space usage. Some examples of this context and its effect on simulation are shown below.

Table 2: How building context shifts employee demand across the validation dataset.

Description	Encl. Focus	Open Focus	Encl. Collab	Open Collab
A downtown finance headquarters.	-5%	+5%	-6%	+5%
An engineering-focused office in the Boston suburbs.	-0%	-5%	+8%	+5%
A biotech campus in the US.	+16%	-7%	+8%	+11%
A regional hub in London, primarily for sales and marketing.	-4%	+1%	+0%	-1%
A headquarters for a technology company in Singapore.	-3%	-8%	+17%	+7%

From Predictions to Confidence: The Simulation Engine

The LSM doesn't produce a single deterministic answer — because real workplace behavior isn't deterministic. Instead, the model predicts *probability distributions* that capture the range of likely outcomes for any given scenario.

What the Model Predicts

For each simulation, the LSM generates three core predictions:

- **Space Popularity:** Which spaces will be used first? For example: *“Conference rooms of capacity 4 are expected to be more popular than those of capacity 8, and will therefore fill up first.”*

• **Gathering Sizes:** How many people will typically be in each space? For example: *“This conference room will often have one person working alone, but about half the time it will host meetings of 4–5 people.”*

• **Employee Demand:** What percentage of people on the floor will need each type of space? For example: *“When the building is busy, about 45% of employees will need a conference room.”*

Turning Probabilities into Actionable Numbers

To translate these probability distributions into concrete, actionable numbers, the LSM runs thousands of Monte Carlo simulations. Each simulation represents one plausible version of how the day could unfold — sampling from every distribution to produce a complete picture of who is where. By running thousands of these simulations, we can make confident, quantified statements such as: *“When 200 people are in this building, there is a 90% chance all employees will find the space they need — but a 10% chance that enclosed collaboration spaces will be oversubscribed.”*

Each simulation follows three simple steps:

1. The **Employee Demand** is used to determine how many of the people will need to use each space type.
2. The **Gathering Sizes** are used to determine how many people each room or space can actually accommodate.
3. The spaces of each space type are filled in **Space Popularity** rank order.

These three simple steps are repeated for every Monte Carlo simulation. At the end of this process, we have “simulated sensor data” (integer person counts) for every single space. This is the final product of the LSM framework.

Validation and Accuracy

How We Validate

The LSM is validated against real-world ground truth — actual sensor data from buildings the model has never seen during training. We hold out a dedicated validation dataset and compare the model’s predicted occupancy patterns against what actually happened. Because the LSM produces probabilistic outputs (ranges of likely outcomes, not single-point predictions), we use *Continuous Ranked Probability Score* (CRPS), a standard statistical metric for scoring probabilistic forecasts. The CRPS value can be interpreted as the **average person-count error per space** — lower is better.

Accuracy at a Glance

Even with no measured data at all, the LSM predicts space usage with an average error of **less than one person per space**. As more context is provided, accuracy improves significantly:

Table 3: LSM accuracy improves as more context is provided.

Context Level	Avg. Error per Space	vs. Baseline
No context (space metadata only)	0.50 people	Baseline
+ Building context	0.50 people	1% improvement
+ Partial measured data (50%)	0.44 people	12% improvement
+ Full measured data (100%)	0.39 people	23% improvement

The takeaway: **the LSM delivers meaningful predictions from day one, and every additional input makes them better.** Organizations can start with just a floor plan and basic space metadata, then layer in building context and measured data over time to continuously improve precision.

Detailed Accuracy by Space Type

Table 4: Average error per space (in people) and improvement from baseline, by space type.

Space Type	No Context	+ Building Context	+ 50% Measured	+ 100% Measured
Enclosed Focus	0.41	0.41 (+1%)	0.34 (+18%)	0.29 (+30%)
Open Focus	0.39	0.39 (+0%)	0.35 (+11%)	0.31 (+21%)
Enclosed Collab	0.79	0.79 (-0%)	0.75 (+6%)	0.66 (+17%)
Open Collab	0.42	0.41 (+1%)	0.35 (+17%)	0.29 (+31%)

Validation Conclusion

The LSM is production-proven and already powering live customer deployments through VergeSense's Predictive Planning product. Across all space types and context levels, the model consistently predicts within less than one person of actual observed occupancy — and accuracy improves substantially as building context and measured data are added. Because VergeSense maintains the ground truth (real sensor data from thousands of deployed buildings), the model is continuously validated and improved against the highest-fidelity occupancy dataset available.

License