Decoding Place Functionality: Identifying Emerging Third Places and the Drivers of Remote Work Visitation via Human Trajectories

Inessa Tregubova*

August 25, 2026

Abstract

The COVID-19 pandemic has reshaped urban dynamics by disrupting the traditional “home-work-third place” triad, giving rise to a new “home-third place-home” spatial pattern. In this structure, 'new third places' (non-traditional work places) increasingly function as office substitutes, moving professional activity from the central business district (CBD) closer to home locations. On the other hand, with the opportunity to work outside the office, home location is moving from the city center to suburbia where housing costs are lower. Existing research acknowledges that among two remote neighborhoods people choose one that provides a bigger diversity of urban amenities. However the specific influence of ‘new third places’ in this choice remains under-researched due to identification complexities. Therefore, this research develops a new methodology to identify emerging “new third places” within the Tel-Aviv Metropolitan Area (TAM) and examine their role in supporting remote work at the neighborhood level. The area choice is explained by the negative migration dynamic in Tel-Aviv between 2020–2022 and retaining work from home (WFH) pattern discovered in our earlier work. To answer the research question, for each place of interest (POI) suitable for remote work we construct the dataset of metrics derived from the GPS-signals collected between January 2020 and September 2023 and extend with features available via Google Places API. The metrics construction involves different techniques starting from the estimation of visiting statistic to measuring the semantic distance of each POI to ‘work’ using Natural Language Processing Word2Vec model. The final identification of office substitutes we produce using the Logistic Regression

*Department of Spatial Science, Hebrew University of Jerusalem

classifier with Elkan and Noto adjustment methodology which allows to have only positive groundtruth. For this we are using POIs with the clear evidence of remote work from Google Reviews.

Finally, to evaluate how this newly identified infrastructure reshapes commercial activity, we employ a stepwise Two-Way Fixed Effects (TWFE) panel regression estimated via Weighted Least Squares (WLS). This spatial econometric framework, widely applied in recent literature to isolate pandemic-induced geographic reorganization, interacts the post-COVID structural shift with distance to the central business district, occupational telework capacity, and local third-place density. Our results reveal that a robust local stock of third places significantly anchors residents and increases localized commercial vitality. However, this magnetic pull is subject to steep spatial decay; while it successfully shrinks travel radii and drives neighborhood consumption in the urban core, it struggles to replicate this retention power in auto-dependent suburban peripheries. Ultimately, these findings provide critical insights into the spatial limits of the “15-minute city” and the uneven geographic footprint of the remote work transition.

1 Introduction

Although the COVID-19 pandemic began six years ago, hybrid and working from home (WFH) patterns are firmly established. The share of remote workers remains steady at roughly 30% in the US (Barrero et al., 2023), 20% to 52% in Europe, and 15% in Israel (Zontag et al., 2022). This stabilization implies a profound spatial reorganization of urban life. As a substantial share of the workforce no longer commutes daily to central business districts (CBDs), their professional, social, and consumption activities have anchored in residential neighborhoods. By shifting daytime populations and demand for amenities toward local areas, the permanent entrenchment of remote work reshapes the very fabric of the city.

1.1 The 15-minute city and expected impacts of remote work

Surging demand for local amenities strongly supports the ”15-minute city” paradigm, which integrates diverse activities so residents can fulfill daily needs locally. Proximity to rich amenities, like neighborhood high streets, is now a defining factor in balancing lifestyle and housing costs (Gokan et al., 2022; Tregubova and Gdaliahu, 2025).

Anchoring remote workers locally yields widely acknowledged benefits, including significant decreases in commuting burdens and carbon emissions (Li et al., 2024). The localized presence of highly educated professionals stimulates suburban development and fosters new low-skilled workplaces (Gokan et al., 2022). However, a limited understanding of exactly which local attractors facilitate this behavior restricts policymakers from proactively adapting city infrastructure to maximize these community benefits.

1.2 The Multidimensional Impacts of Localized Remote Work

Beyond real estate changes, the literature highlights diverse impacts introduced by the rise of informal third places and collaborative neighborhood spaces.

Economic Revitalization and Labor Market Shifts Teleworking reallocates local consumption services (LCS) spending from CBDs to residential areas. This offsets CBD job losses by creating local service jobs, raising wages for unskilled workers, and reducing the skilled-unskilled wage gap. Third places also draw sustained foot traffic into peripheral communities, revitalizing commercial activities while offering corporate employers 20% to 40% cost efficiencies through ”hub-and-spoke” real estate models.

Urban Regeneration and Physical Space Transformation Third places serve as catalysts for physical regeneration by adaptively reusing vacant infrastructure, such as abandoned factories or schools. Retrofitting these spaces combats urban decay and spurs microscale physical improvements, like temporary cultural installations and newly equipped community leisure spaces.

Environmental Sustainability and Green Mobility The local shift mitigates vehicle miles traveled (VMT), alleviating traffic congestion, energy consumption, and greenhouse gas emissions. Local proximity naturally encourages remote workers to adopt active transportation modes like walking or cycling.

Community Well-being, Networking, and Social Innovation Local third places help remote workers bypass household distractions and elevated energy consumption. By providing a collaborative atmosphere, they mitigate social isolation, boost job satisfaction, and facilitate knowledge spillovers among independent professionals. Socially oriented third places act as collective community goods (Mariotti et al., 2023), promoting inclusion via digital training or functioning as emergency hubs during crises. Ultimately, these spaces create vital ecosystems in peripheral regions that retain local talent and prevent ”brain drain.”

1.3 The shift in local activity and the rise of new third places

A critical component of this local shift is the spatial distribution of work activities. Adapting Oldenburg (1989)’s classic definition, ”new third places”—such as coworking spaces, local cafés, libraries, and community centers—attract remote workers by providing reliable Wi-Fi, power outlets, and a comfortable environment (Li et al., 2024; Tomaz and Tabrizi, 2024; Yang et al., 2025). These hybrid spaces serve as vital office replacements, fulfilling the need for ambient social interaction (Tomaz and Tabrizi, 2024). Their success relies heavily on high public transport accessibility and proximity to remote workers’ homes (Li et al., 2024).

Despite their growing visibility, existing literature lacks robust quantitative proof regarding the usage intensity of these places compared to working directly from home. Furthermore, no research has thoroughly examined the extent to which specific third places (beyond traditional coworking spaces) explain local remote working levels and neighborhood housing price dynamics. Previous studies primarily capture correlational associations and fail to adequately address potential endogeneity problems (Yang et al., 2025).

This study addresses this gap by identifying ”new third places,” mapping their utilization, and measuring how this new work-leisure function affects local housing prices. We detect existing places adopting this new use by applying recently developed approaches for measuring human behavior and Point of Interest (POI) functionality. Behavior is measured using trajectories from GPS-signal datasets, which capture detailed spatial patterns and hidden connections invisible in official statistics and surveys.

To mitigate inherent biases in mobile trajectory data—such as disproportionate demographic representation or varying app usage—we validate our findings against Google Maps Reviews as a ground truth sample. Recent research by Li et al. (2026) demonstrates this data source's efficacy in capturing the real, lived functionality of POIs, which is often impossible to identify solely by observing physical characteristics.

The paper proceeds as follows: first, we introduce the Tel Aviv Metropolitan (TAM) study area and explain the need for this research; second, we describe the data and methodology; third, we present and discuss empirical results; and finally, we conclude with our contribution to the broader urban economics literature.

2 Description of the Data

2.1 GPS signals dataset

Empirical testing of the model requires a method for measuring the dynamics of remote work at the neighborhood level before and after the COVID-19 pandemic. To this end, we utilize a dataset of GPS signals collected via mobile applications in TAM between January 2019 and September 2023. The dataset contains geolocated 'stays' - clusters of consecutive signals occurring at the same location. For each stay we have User ID, the beginning and end of the stay and location centroid. After removing users with less than 4 'stays' per month, the data capture approximately 1% of the TAM population before 05/2020 and 5% after on a monthly basis (Figure 1). The daily hours distribution is more stable but still reveals the outliers in the month. Such differences in data volume and signal variation significantly limit our ability to compare pre- and post-COVID periods in the next steps and require us to normalize signals volumes before drawing conclusions about temporal differences in hourly patterns.

Figure 1: Monthly temporal statistics of the GPS signals dataset

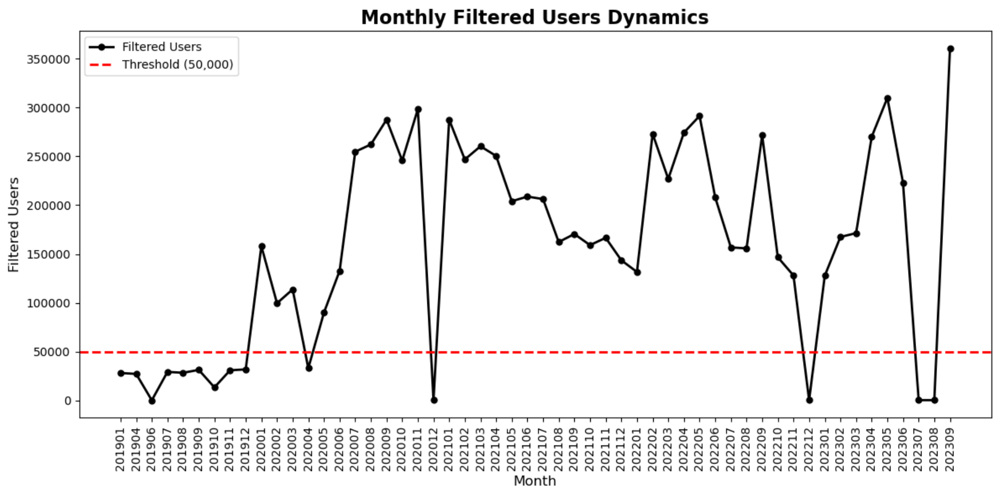

Notes: The top panel shows number of users with at least 4 stays in a month; the bottom panel shows the hourly distribution of stays across the period. Both diagrams show the significant variations in the stay counts suggesting the removal of some months and normalization from further analysis.

2.2 POIS Datasets

To map signals to physical locations, we utilize several sources for Points of Interest (POIs), specifically the Google Places API and OpenStreetMap (OSM). Although OSM is a frequent choice in academic research, it exhibits a high degree of spatial discrepancy that can potentially compromise mapping quality. Consequently, to ensure accuracy in identifying third places, we collected POIs from both sources and then select the best candidate for signal location during the later steps. First of all, we extracted 4902 potential third places within the TAM from Google Places for the following keywords: *bakery, cafe, cafeteria, coffee_shop, tea_house, dessert_restaurant, cat_cafe, dog_cafe, breakfast_restaurant, brunch_restaurant, library, book_store, museum, cultural_center, community_center.* From them we keep 4176 places that have at least one review on Google Maps. Next, we collected POIs indicating work or education locations such as *business_center*, *industrial_area, corporate_office* (Google Places) and *university, school* (OSM). Lastly, we collected data on the rest urban and natural amenities from OSM to detect these stops in users daily trajectories. The full list can be find in the Appendix A.

For each place int Google Places dataset, we retain a unique Google identifier, place name, Google primary place type, location, ranking, the number of reviews, and the data on availability of specific amenities - such as Wi-Fi or restroom. Additionally, we collect five 'most relevant' reviews from Google Maps – the maximum number available through service

API. The limitation of the data.

2.3 Statistical data

Neighborhood controls come from the Israel Central Bureau of Statistics (CBS) at the statistical-area (SA) level. From 2019 we use population totals, the share aged 20–59, and the CBS socio-economic cluster, mapped from 2011 SA codes to 2022 boundaries with population weights. From the 2022 Census we add SA industry and occupation shares on the same geography (including IT, finance, academic occupations, and academic degrees). TAM contains 1473 SAs, in which 1241 are inhabited.

3 Strategy for identifying third places

Since the goal of the research is to estimate the impact of third places on housing prices, the methodology includes two steps: identification of third places and estimation of their impact on housing prices in the period 2021-2023.

The identification of third places is the complex task that in the light of uncomplete information in available sources, requires to embrace multiple approaches and sources, and validate them versus each other.

For this matter, the identification strategy contains several steps:

- Defining a sample of candidate office substitutes
- Identifying the sample of ground truth third places based on Google reviews
- Building the identification model using Positive-Unlabeled (PU) learning
- Conducting a descriptive analysis of third places versus alternative venues based on GPS signal data
- Analyzing the role of third places as office substitutes within users' weekday trajectories using location embeddings

3.1 Defining a sample of candidate office substitutes

The sample that we use to find third places candidates is limited by POIs collected from the Google Places database, which were selected based on the list of keywords presented in the previous chapter.

Next, we restrict the sample to residential and mixed-use areas, excluding formal workplaces and institutional sites. This approach aligns with existing literature indicating that third places are primarily located near homes outside of commercial hubs. Furthermore, removing such places prevents the potential misclassification of office workers' daytime visits as remote work from the third place. For this matter, we construct the *location_fits_remote_work* flag that indicates if the surrounding building environment, official land-use, or local POI composition refers to a formal work location.

To identify work zones, we spatially intersect these locations with datasets of dedicated employment zones and Census Statistical Areas (SAs). A location is designated as a work zone if it falls within an "industrial" land-use area according to the Census or an employment zone entirely lacking residential buildings. Furthermore, if a geohash intersects with polygons representing universities, colleges, medical institutions, or transport hubs, it is classified as an institutional zone. Lastly, we calculate the number of business centers and corporate offices within a 250-meter radius of the geohash; if this count exceeds nine (falling into the top decile), the area is considered a work zone. Areas containing one to nine of these establishments that are not otherwise classified as work or institutional zones are categorized as mixed-use. Finally, all locations classified as work zones or institutional are marked as unsuitable for remote work (*location_fits_remote_work* = False). The remaining locations are considered candidate third places. Moving forward, we analyze only POIs with *location_fits_remote_work* = True.

3.2 Identifying the sample of ground truth third places based on Google reviews

To identify third places, we rely exclusively on user reviews and amenity features from Google Maps. When individuals mention working on laptops or use synonymous terms in a review, it provides a strong indication of remote work activity. A limitation of the Google Places API is that it returns a maximum of five "most relevant" reviews. Consequently, we are likely to miss reviews indicating remote work for many locations, resulting in potential under-identification (low recall). However, the accuracy rate for places that *do* possess such reviews is expected to be high (high precision), allowing us to utilize this subset as a ground-truth positive sample for our classification model.

To operationalize this, we derive a text-based indicator *flag_keyword_remote_work*. *Direct remote-work evidence* is identified using a dictionary of explicit work-related terms and phrases, such as "laptop" or "place to work" (the full list is presented in Appendix B). A review contributes positive evidence only if it contains at least one such term and does not contain exclusion terms that indicate unrelated contexts (e.g., ”event venue” or ”no Wi-Fi”; see Appendix B). A location is flagged as positive when at least one qualifying review is detected.

To bolster this dataset and ensure a robust sample, we augment the initial 100 ground truth locations with venues possessing a suitable working atmosphere. We construct *flag_atmoshere_remote_work*, leveraging keywords that reflect a calm social ambiance (e.g., ”quiet”, ”meeting”, ”sitting with others”) while explicitly excluding negative ambiance signals (e.g., ”loud”, ”rush”). Locations whose reviews indicate this suitable atmosphere and that serve coffee are added to the ground truth positive sample. Finally, we extend the feature set with amenity availability parameters returned via the Google Places API, though these are filled by owners and missing for approximately 50% of places.

3.3 Identification of third places

Based on the dataset of candidate third places and their extracted Google Reviews features, we formulate the identification task as a Positive-Unlabeled (PU) learning problem. This binary classification approach is suited for our data because we lack verified negative samples; the absence of a remote-work review does not definitively prove a location is not used for remote work, it simply means we did not observe the review.

Following the Google API description, we assume the randomness of review selection, meaning the chances of retrieving a relevant review for any true third place are equal. This allows us to utilize the adjustment methodology proposed by Elkan and Noto (2008), where the classifier estimates the probability that an observation is labeled positive, *P(s=1|x)*, and assumes this is proportional to the target probability *P(y=1|x)* by a constant *c*. The final probability is computed as: $P(y = 1|x) = P(s = 1|x)/c$.

Given the potential multicollinearity among the text-based and amenity features, we utilize an Elastic-Net Logistic Regression model to maintain coefficient stability. To interpret feature strength derived from the coefficients, we standardize the inputs using the StandardScaler method.

Since ground-truth negative examples are unavailable, we evaluate the model using a repeated positive-holdout procedure. In each repetition, we randomly hide 30% of the trusted positive places (changing their label to unlabeled), train the Elkan-Noto PU model on the remaining data, and rank all candidates by the predicted probability. Model quality is assessed by whether the hidden positives are recovered near the top of this ranking, using *recall-at-k* measures and median percentile ranks (Saunders & Freitas, 2025). We prioritize minimizing False Negatives to identify as many third places as possible, accepting a small margin of 'wrong' third places. Additionally, we test model sensitivity across hyperparameters, iterating the regularization coefficient C over [0.03, 0.1, 0.3, 1] and the *l1-ratio* over [0.25, 0.5, 0.75].

3.4 Collecting POI visitation metrics using GPS signal data

The goal of this stage is to extract weekday visitation metrics for each place from Google places dataset based on GPS signals to conduct a descriptive analysis of new residential behaviors. By deriving these metrics, we aim to empirically test hypotheses from the literature by evaluating whether identified third places exhibit distinct utilization patterns—specifically, functioning more akin to formal office substitutes—when compared to other POIs within the same amenity category.

The process begins by classifying each *stay* as Home/Work/POI/Unknown. This involves a hierarchical spatial assignment: GPS signals are first matched against the user's Home or Work Geohash Spatial Index, and subsequently against a 20-meter POI radius if the initial match fails. Home and work locations are recalculated monthly using our established frequency-based methodology (Tregubova & Gdaliahu, 2025). This specific radius was selected to balance the inherent inaccuracy of GPS positioning against the risk of misattributing signals to incorrect, neighboring POIs.

Matching the January 2019 to January 2020 signals to individual POIs revealed a very low average visit volume of fewer than 10 (Figure 2). This reduction is a direct consequence of the fivefold difference in total stays between this period and subsequent years identified earlier. Therefore, for our descriptive analysis of shifting mobility patterns, we utilize only the post-COVID period data.

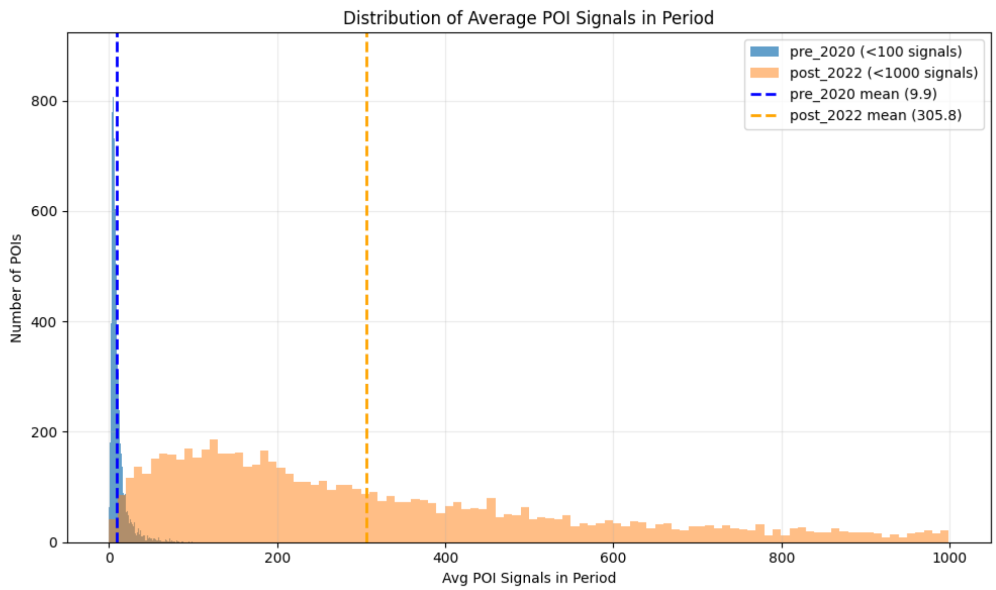

Figure 2: Distributions of average visits per individual POI in two periods: 01/2019–01/2020 and 01/2022–09/2023.

Following this spatial attribution, we filter the dataset to keep only POIs with a minimum of ten visits during 01/2022–09/2023. To test the hypothesis that remote workers exhibit distinct temporal footprints in third places (Tomaz and Tabrizi, 2024), we calculate the *average share of remote work visits*. We define a "remote work visit" as a continuous stay lasting more than one hour that occurs between 10:00 AM and 3:00 PM. This classification aligns with documented behavioral patterns suggesting that remote workers avoid peak commuting hours, arrive at third places between the morning rush and the lunch peak when venues are less crowded, stay for longer durations, and typically originate from a residential location.

To further test the hypothesis that alternative workspaces exhibit lower visitation variance than locations driven by discrete consumption events (e.g., meal or coffee rushes), we construct metrics representing visitation stability relative to a POI's daily peak: *hourly signal standard deviation* and the *no_drops_in_daily_dynamic* flag. To mitigate the uneven hourly sampling inherent in GPS data collection, we first normalize the signal counts at a given POI against the total signal volume recorded during that corresponding hour. The first metric calculates the standard deviation of these normalized hourly averages per POI. The second sets a binary flag to 1 if the minimum visit volume between 10:00 AM and 3:00 PM remains at or above 20% of the peak hour volume. These metrics empirically capture the stabilizing effect of remote workers on daytime foot traffic.

Beyond these targeted behavioral indicators, we extract a comprehensive list of POI-level features to test broader spatial and temporal hypotheses about localized visit intensity. We compute the average *visit frequency* per unique user, the *absolute volume of weekday visits*, and the average *dwell duration* on weekdays. We also calculate the overall share of visits occurring on weekdays versus Fridays to capture overarching temporal shifts.

Finally, to assess the hypothesis that remote work drives the localization of amenities and their integration into immediate residential neighborhoods, we examine the spatial origins of visits. For each visit, we identify the user's previous location, as well as their established home and work locations. From this, we derive the distance to the user's formal work and home locations, and an indicator of whether the user originated from home. We also calculate the proportion of visits occurring within a one-kilometer radius of the user's home. Together, these temporal and spatial features yield a multidimensional dataset of 11 metrics specifically designed to describe and validate new mobility patterns surrounding urban amenities.

3.5 Detecting work replacers in users weekday trajectories using location embeddings

The final indicator is the most comprehensive - it reveals not only a specific location's potential to replace a formal office but also broader shifts in resident behavior. For this analysis, we utilize the previously generated dataset of anchored stays, where each instance is classified as home, work, POI, or unknown.

Detecting office substitutes requires constructing location embeddings: vector representations of each place derived from its sequential position within daily mobility trajectories (Esuli et al., 2018; Murray et al., 2023; Weng et al., 2025). Locations sharing similar trajectory contexts—meaning they are frequently preceded and followed by the same spatial sequences—yield mathematically proximate embeddings. As such, third places that play role of office substitutes should have embeddings similar to established work locations. we measure similarity using cosine distance (Crivellari and Resch, 2022).

We construct embeddings in three steps: first we create vocabulary of 'tokens' – descriptions of each places, then we build dataset of individual mobility trajectories, and last we train on them a Word2Vec Natural Language Processing (NLP) model (Mikolov et al., 2013) to create location embeddings.

Each location is encoded into a discrete 'token' two times based on two groups of encoding rules: one for aggregated naming and one for individual naming (Table 1). The first group of rules is applied to both pre- and post-COVID periods and is used for descriptive purposes: to identify changes in users behavior. The second dictionary of tokens we apply only to post-COVID data in order to calculate cosine distance between individual POI location in different time of the day and ‘work’.

For each vocabulary of ‘tokens’ we create a dataset of trajectories. Each trajectory represents chronologically ordered sequences of places (‘tokens’) visited by a user within a specific day. To ensure the robustness of this movement representation, we enforce a minimum stay duration of 10 minutes and a maximum temporal gap of three hours between visits. If the time gap exceeds three hours, the trajectory is terminated, and a new one is initiated. In the end, we retain only trajectories containing a minimum of three locations.

	Home or Work	**POI**	**Unknown location**
Aggregated locations	‘Home’/‘Work’	‘POI’ + [POI category]	‘GH’ + [Day Period]
Individual locations	‘Home’/‘Work’ + [Statistical Area ID]	‘POI’ + [Unique ID] + [Day Period]	[Geohash Index] + [Day Period]

Table 1: Encoding rules for different types of locations

Following this procedure, we yield a final dataset of 5,744,898 ‘sentences’: 855,538 in pre-COVID period and 4,889,360 in Post-Covid. Aggregated version of tokens comprises 42 unique location tokens with an average frequency of 25,972 appearances in Pre-Covid period and 134,574 in Post-Covid period. While for the size of tokens with unique identifiers is 57,503 with an average frequency of 99.8 appearances in post-covid period.

Leveraging these discrete trajectories, we train a Word2Vec model using the Gensim Python library. Guided by literature best practices and scaled to the size of our dataset, we apply the following model parameters:

- Architecture: skip-gram
- Embedding dimensionality: 100
- Minimum token frequency: 5
- Negative sampling: 15 negative samples per positive pair
- Training epochs (passes over the corpus): 10

We execute the model 3 times: two times for aggregated tokens in two periods and one with individual tokens for post-Covid period. Then, for each model separately, we identify POIs that are semantically closest to ”work” locations by calculating their cosine similarity. This metric mathematically translates associative strength into the angle between normalized vectors within the embedding space. Last, to account for the average differences in similarity to ‘work’ between POIs categories, we replace original cosine similarity of the place with its ranking inside the category.

The obtained results are used to visualize the token embeddings from each model in T-SNE representation and measure distance to ’WORK’ and identify the most frequent 3-grams (three consecutively visited locations) to descriptively illustrate the differences between two periods and conclude about emerging patterns in post-COVID era.

Empirical Strategy & Two-Way Fixed Effects Framework

To evaluate the shifting spatial equilibrium of out-of-home commercial consumption, we employ a stepwise Two-Way Fixed Effects (TWFE) panel regression model estimated via Weighted Least Squares (WLS). The models progressively interact the post-pandemic structural shift with spatial gradients (distance to the core), neighborhood telework capacity, and local amenity availability. The unit of observation is Census 2022 Statistical Area (SA)

The baseline specification leverages the TWFE structure to control for unobserved, time-invariant neighborhood characteristics as well as metropolitan-wide temporal shocks. Because of this strict two-way fixed setup, the constituent un-interacted terms (Distance, WFH Potential, and Amenity Dose) are perfectly collinear with the spatial fixed effects and are consequently absorbed. This allows the models to cleanly isolate the differential effects of the pandemic disruption across the spatial hierarchy. To account for heteroskedasticity across SA of varying population sizes, the models are weighted by the pre-pandemic (January 2020) *worker* count within each SA, with standard errors clustered at the SA level. *Worker* in this context is an individual in the GPS-signals dataset for whom we identified first (home) and second (work) frequent place.

Variable Construction

To parameterize the spatial and built-environment heterogeneity, we construct the following cross-sectional features:

- **Distance Gradient ($Dist_i$):** The physical distance, measured in kilometers, from the home Statistical Area (SA) to the Central Business District (anchored at Azrieli Center, Tel-Aviv). To account for non-linear spatial decay, it enters the models log-transformed as $\ln(1 + Dist_i)$.
- **Telecommuting Capacity (WFH_POT_i):** A localized proxy for occupational remote work capacity, utilizing the Information Technology (IT) employment share

within the SA, sourced from the 2022 Central Bureau of Statistics (CBS). Missing values in the census data are imputed as 0.

- **Amenity Dose ($Dose_i$):** A time-invariant metric capturing the local stock of commercial venues. For each worker active in the baseline period (January 2020), we calculate the count of venues within a 1 km radius of their home geohash with at least 5 visits within a month, which is then averaged at the home SA level. To mitigate right-skewness, the raw SA averages are transformed using $\ln(1+x)$. The models iteratively evaluate three dose variations:
  - *Log_Cand*: The log-transformed count of candidate venues.
  - *Log_TP*: The Log-transformed count of positively identified third places.
  - *TP_Ratio*: The localized saturation of third places, computed as $TP_n/(Cand_n+1)$.

To systematically evaluate how local Third Places reshape commercial activity post-COVID, we test several hypotheses by comparing the pre-pandemic baseline (January and February 2020) with the post-disruption period (Q1 2023). We define the "local area" as a 1 km buffer around each mobile user's home location. The models evaluate five distinct dependent variables (Y_{it}), aggregated at the SA-month level:

- **Substitution of "global" POIs for local ones:** The share of local POI hours relative to all POI hours.
- **Shift of local time from home to amenities:** The share of local POI hours relative to all local hours.
- **Preference for local third-place candidates:** The share of local candidate hours relative to all candidate hours (where candidates are defined as Google Places in categories commonly used for remote work).
- **Overall affinity for out-of-home activities:** The share of total POI hours relative to all hours.
- **Amenities travel radius:** The Radius of Gyration specifically for POI visits, calculated in meters as the root-mean-square (RMS) distance from the center of the users' activity. This operationalization follows standard spatial mobility frameworks frequently applied in urban analytics research.

Stepwise Model Specifications

Let Y_{it} represent the mobility or consumption outcome (e.g., local POI retention share, radius of gyration) for Statistical Area i at time t. The progression of the four regression models is formalized as follows:

M1 (Distance Gradient): Establishes the baseline post-pandemic spatial decentralization effect.

$$Y_{it} = \beta_1(Post_t \times \ln(1 + Dist_i)) + \gamma_i + \delta_t + \epsilon_{it}$$

M2 (Telecommuting Capacity): Introduces local work-from-home potential and its spatial interaction to isolate the independent effect of localized remote work adoption.

$$\begin{aligned} Y_{it} = {} & \beta_1(Post_t \times \ln(1 + Dist_i)) + \beta_2(Post_t \times WFH\_POT_i) \\ & + \beta_3(Post_t \times \ln(1 + Dist_i) \times WFH\_POT_i) + \gamma_i + \delta_t + \epsilon_{it} \end{aligned} \tag{1}$$

M3 (Amenity 1km Buffer): Excludes WFH potential to independently test how local amenity density (Candidate venues or Third Places) draws consumption into the neighborhood post-pandemic.

$$Y_{it} = \beta_1(Post_t \times \ln(1+Dist_i)) + \beta_4(Post_t \times Dose_i) + \beta_5(Post_t \times \ln(1+Dist_i) \times Dose_i) + \gamma_i + \delta_t + \epsilon_{it}$$

M4 (Full Model): The saturated specification evaluating the simultaneous and interactive spatial effects of both remote work potential and local amenity infrastructure.

$$\begin{aligned} Y_{it} = {} & \beta_1(Post_t \times \ln(1 + Dist_i)) + \beta_2(Post_t \times WFH\_POT_i) \\ & + \beta_3(Post_t \times \ln(1 + Dist_i) \times WFH\_POT_i) + \beta_4(Post_t \times Dose_i) \\ & + \beta_5(Post_t \times \ln(1 + Dist_i) \times Dose_i) + \gamma_i + \delta_t + \epsilon_{it} \end{aligned} \tag{2}$$

Additional Terms Defined:

- $Post_t$: A binary indicator taking the value of 1 for the post-disruption observation period (2023Q1), and 0 otherwise.
- γ_i: Statistical Area (SA) unit fixed effects.
- δ_t: Time (year-quarter) fixed effects.
- ϵ_{it}: The robust standard error term.

4 Results

4.1 Google Places sample

In the original sample of candidate work substitutes from Google Places with at least one review, we identified 662 located in work zones, inside institutions or big transport hubs. They were marked as impossible for remote work and excluded from the model sample. Therefore, the final sample contains 3524 candidates.

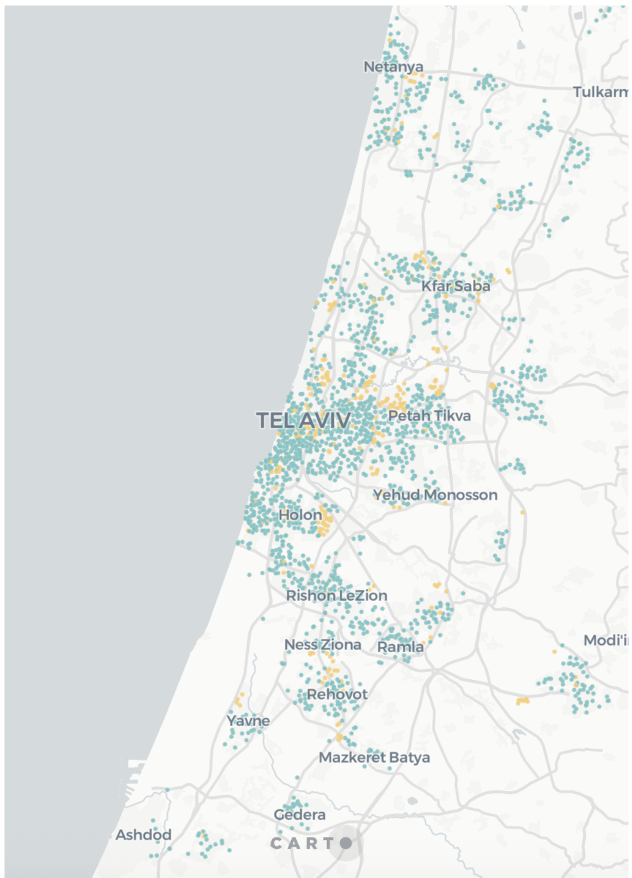

Figure 3: Spatial distribution of third place candidates in TAM, categorized by office and residential locations

In the sample we found 100 places with direct keywords from the reviews indicating remote work. The biggest number of ground truth third places belong to coffee shops (41 places) and cafes (34), then libraries - 12 places, cultural centers and bakeries – 6 places each, community centers – 4 and 1 bookshop. Keywords showing the right atmosphere for remote work were found in 532 places.

4.2 Model results

Evaluating the model over 100 runs with 30% of the positives hidden and C=0.1 and l1 ratio=0.5 revealed an average median rank in the top 11%, with a worst-case performance bounded within the top 18% (Figure 4).

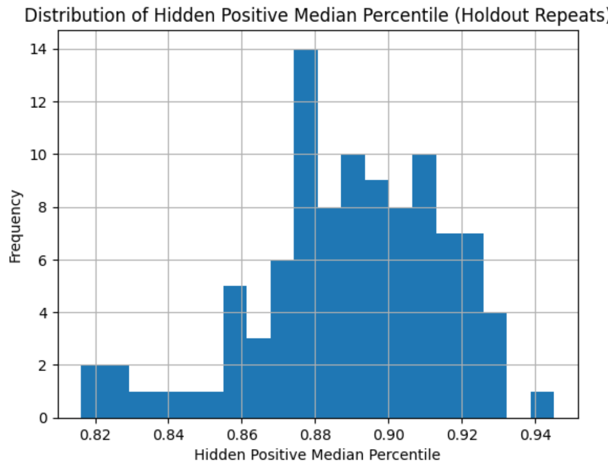

Figure 4: Distribution of Hidden Positive Median Percentile on 100 iterations with 30% of the hidden positive.

By averaging the coefficients obtained at each iteration, we reveal the model's strongest features (absolute coefficient value > 0.1). The results are presented in Figure **??**. The strongest features come from Google Places, with keywords indicating atmosphere being the most prominent. Important to note, that all of them were different from zeros at least in 95% of iterations.

Then, the iteration over the grid of hyperparameters helped us find the best combination, while generally demonstrating that the results are not highly sensitive to their choice. As such, ranking differences exist in the top 20% to 40% range and are fully eliminated by the top 50%, where all specifications demonstrate a 95% recall level. The optimal model configuration for the top 20% of the ranking relies on stronger regularization with an equal contribution from Ridge and Lasso penalties (C=0.03, L1 ratio=0.5). At the top 20%, it reaches a recall of 77%. This configuration we use to produce the final identification.

4.3 Detected 3rd places

The model detected 550 third places, which represents 15.6% of the total candidate sample. The highest shares are observed in the central Tel-Aviv area—22% in the CBD and 19.7% in residential areas—while in satellite cities, we identified 12% of candidates as real third places, regardless of how active the neighborhood is. Therefore, these results suggest that

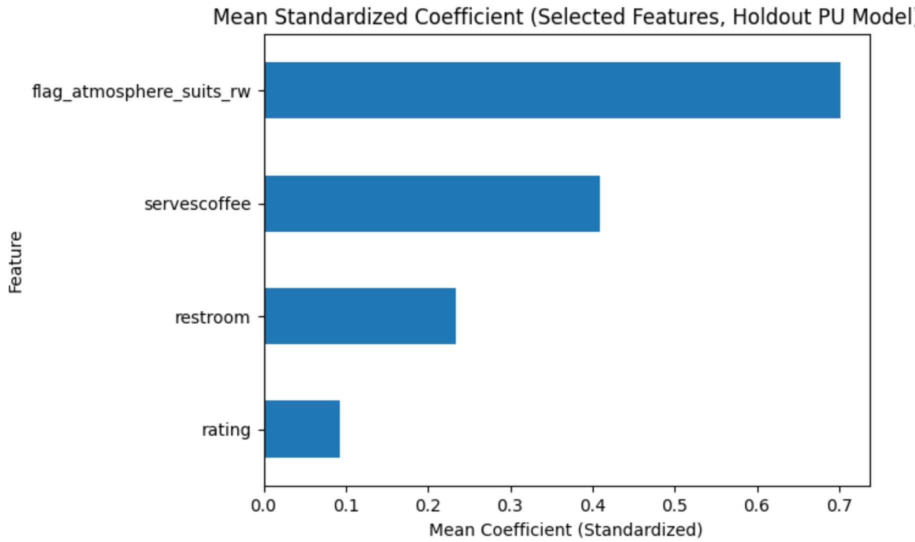

Figure 5: Average feature importance on 100 repetitions with 30% of hidden groundtruth

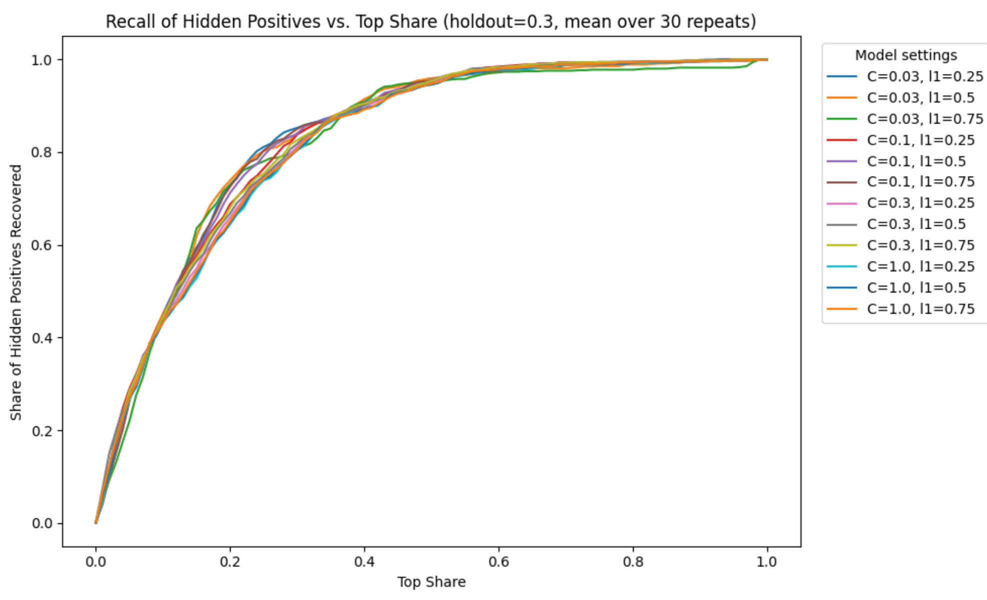

Figure 6: Recall of hidden positives by top share ranking across evaluated hyperparameter configurations (mean over 30 repeats, 0.3 holdout). Each line represents the results for unique combination of hyperparameters. Model performance exhibits low sensitivity to parameter variations, with all specifications converging to 95% recall at the top 50% threshold.

in general the chances of finding work substitutes are higher in central areas rather than suburbia in TAM.

Categorically, we observe that the absolute values are highest among coffee shops (174 places) and cafes (128 places), while the share of places suitable for remote work is highest for libraries.

Primary type	Total places	Model identified places	% of total	Groundtruth count	Groundtruth identified by model	Recall
coffee_shop	581	174	29.95%	41	29	70.73%
cafe	621	128	20.61%	34	17	50.00%
bakery	1115	66	5.92%	6	2	33.33%
cultural_center	392	62	15.82%	6	6	100.00%
library	145	50	34.48%	12	12	100.00%
community_center	243	27	11.11%	0	0	NaN
book_store	246	20	8.13%	1	1	100.00%
museum	126	15	11.90%	0	0	NaN
breakfast_restaurant	11	3	27.27%	0	0	NaN
dessert_restaurant	15	2	13.33%	0	0	NaN
tea_house	23	2	8.70%	0	0	NaN
dog_cafe	3	1	33.33%	0	0	NaN
brunch_restaurant	2	0	0.00%	0	0	NaN

Table 2: Model performance by primary POI type

We also note that the model detects remote work behavior in categories absent from the ground truth sample, such as museums, community centers, and specific niche establishments like breakfast and brunch restaurants, dessert shops, and dog cafes. The model exhibits the highest level of noise within the bakery category, where only 2 of the 6 locations from the ground truth were correctly identified. Spatially, while the absolute number of identified third places is highest in Tel-Aviv, the highest proportional shares are found in satellite cities such as Kfar Saba, Raanana, Kiryat Ono, and Rehovot.

According to the distribution of identified third places by POI category and neighborhood type (Figure 7), coffee shops are the most popular option for remote work across all neighborhoods in absolute terms. Surprisingly, in relative terms the highest share of coffee shops utilized for remote work is found in central residential areas not in the CBD (36.2% vs. 33.5%), where one might typically expect the greatest concentration of such venues. residents of quiet suburbs choose libraries and cultural centers much more frequently than residents of other areas (35 locations versus 10 elsewhere), while the percentage of utilized libraries is highest in the centers of satellite cities.

4.4 Visits distribution analysis

The descriptive analysis of the metrics collected based on GPS-signals show their potential to distinguish third places. Figure 8 illustrates the hourly dynamics of long visits by comparing

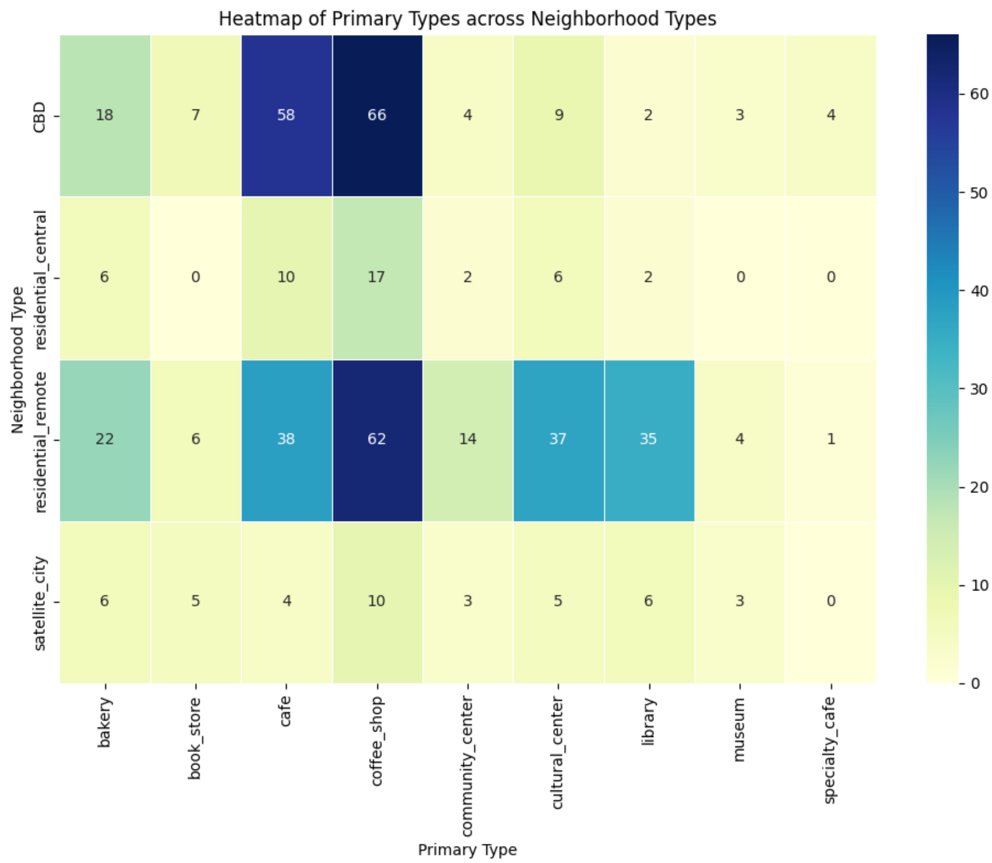

Figure 7: Distribution of identified third places (absolute counts) and Distribution of identified third places (percentages)

two Tel Aviv cafes: a bakery lacking remote work amenities (left panel) and a spacious courtyard café highly conducive to laptop work (right panel). While both reside in mixed-use areas near—but not directly adjacent to—established workplaces, their visitation patterns diverge significantly. In the left diagram, the absolute number of long visits remains stable—accounting for less than 30% of total visits until the 5:00 PM end of the workday—with a relative peak occurring during the 12:00 PM to 1:00 PM lunch hour. Conversely, the right diagram exhibits a marked increase in long visit counts between 10:00 AM and 3:00 PM, culminating at 2:00 PM when over half of all visits exceed one hour.

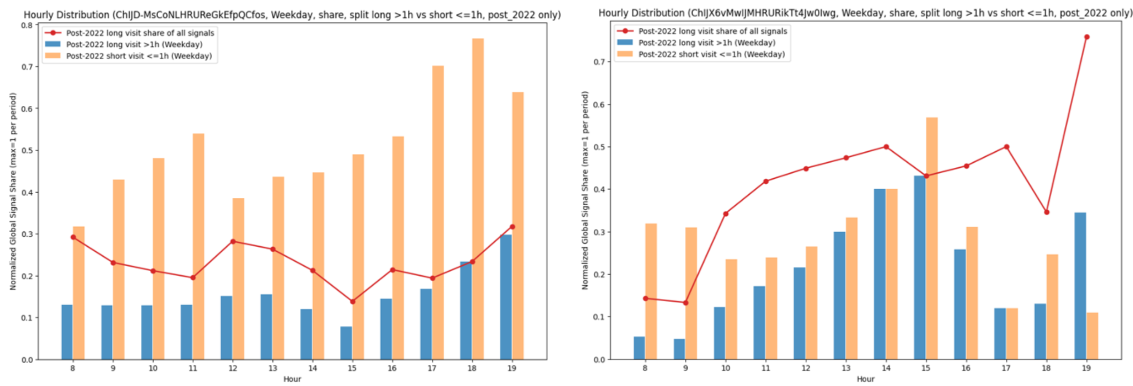

Figure 8: Comparison of normalized visit counts between two locations. A visit is attributed to a specific hour if it overlaps with that hour (i.e., started before or during, and ended during or after). Presented values are normalized against total signals at specific hour and the maximum hourly visit count recorded within the 01/2022–09/2023 study period at specific location.

The analysis of standardized behavior differences reveals that detected third places exhibit distinct visitation signatures compared to other candidate locations (Figure 9). Across the entire Tel Aviv Metropolitan area, the most consistent defining feature of a third place is significantly longer visit durations, averaging an increase of 0.15 standard deviations. However, specific behavioral patterns diverge depending on the geographic context. In the urban core of Tel Aviv-Yafo, third places are characterized by notably fewer arrivals directly from the workplace, showing a decrease of 0.15 standard deviations compared to standard commercial venues. Conversely, in the satellite cities comprising the rest of Gush Dan, the primary behavioral deviation for identified third places is a higher prevalence of extended early-afternoon stays, registering an increase of 0.17 standard deviations.

While the broader hourly distributions largely follow similar temporal rhythms between detected third places and other candidates, a closer examination reveals localized temporal shifts (Figure 10). Notably, third places attract a distinctly higher share of highly localized visits—defined as those originating within 1 km of the user's home—specifically during the

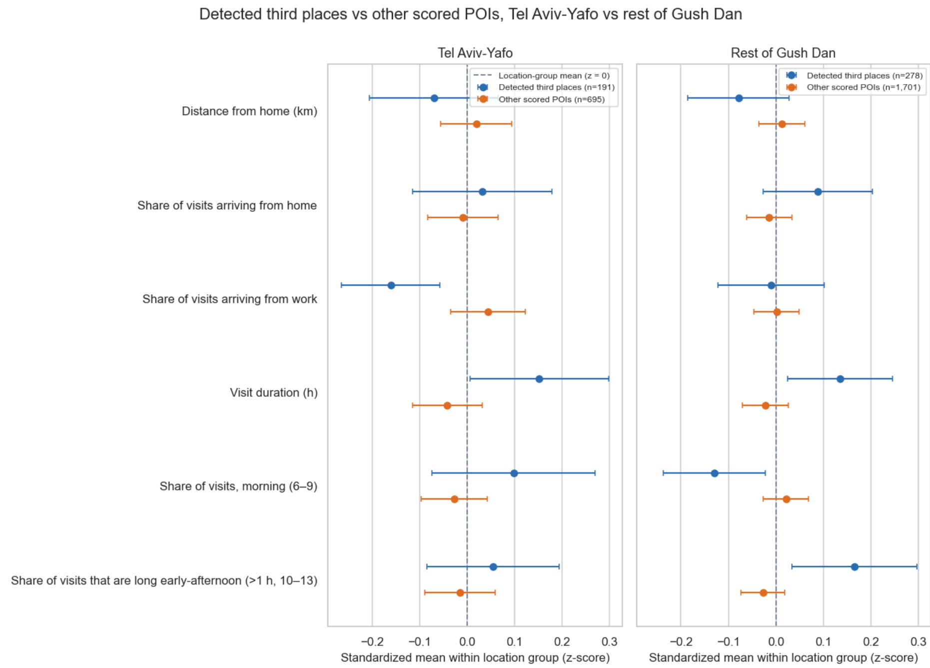

Figure 9: Place-level weekday visit behaviour of PU-detected third places versus the remaining scored POIs, split by Tel Aviv-Yafo versus the rest of TAM. Features are z-scored independently inside each location group. Points are group means; error bars are 95% normal confidence intervals. Unit of analysis is the place; weekday timing shares are among weekday visits 2022-2023, morning hours: 6–9, early afternoon: 10–13

morning hours. This localized morning peak is particularly evident in commercial locations within Tel Aviv between 8:00 AM and 10:00 AM, as well as in civil locations across the rest of Gush Dan from 8:00 AM until noon. Furthermore, supporting the aggregate duration findings, commercial third places outside of Tel Aviv experience a targeted increase in the share of visits lasting longer than one hour during the late morning and early afternoon.

4.5 Trajectories analysis

The analysis of daily trajectories constructed based on GPS signals revealed several shifts in users behavior after COVID-19 that may be indication of remote work presence.

First of all, top 3-grams constructed from aggregated token reveal two new popular trajectories in post-COVID time: Home->Home->Home (2.4% of trajectories) and Home->Park during weekday late morning -> Home (0.7% of trajectories) (Figure 11).

Additionally, we notice that residents of CBD area (within 5 km from Tel Aviv Azrieli center) have started to show a unique new behavior pattern: Home -> food place during late morning hours - > Home (0.27%) (Figure 12). This may signals for remote work from the third places.

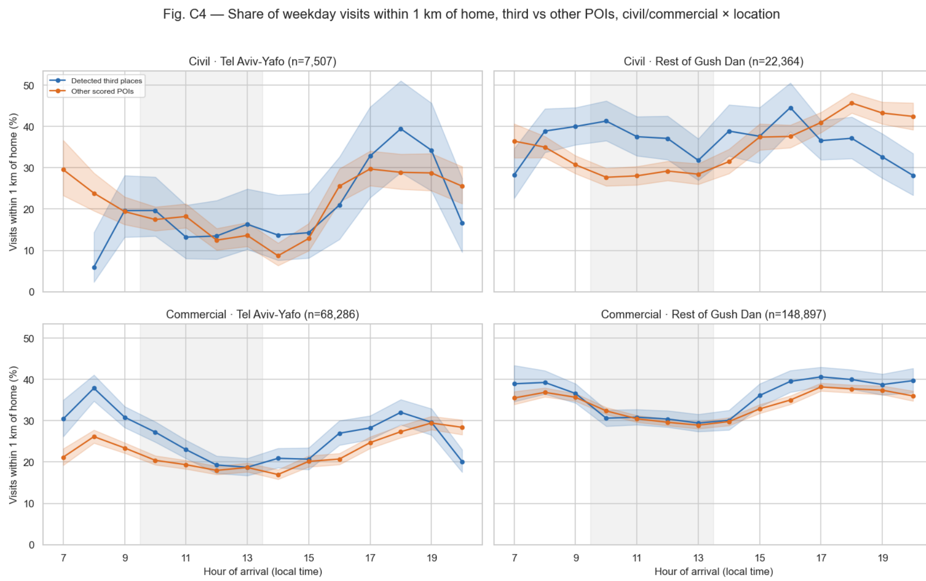

Figure 10: Enter Caption

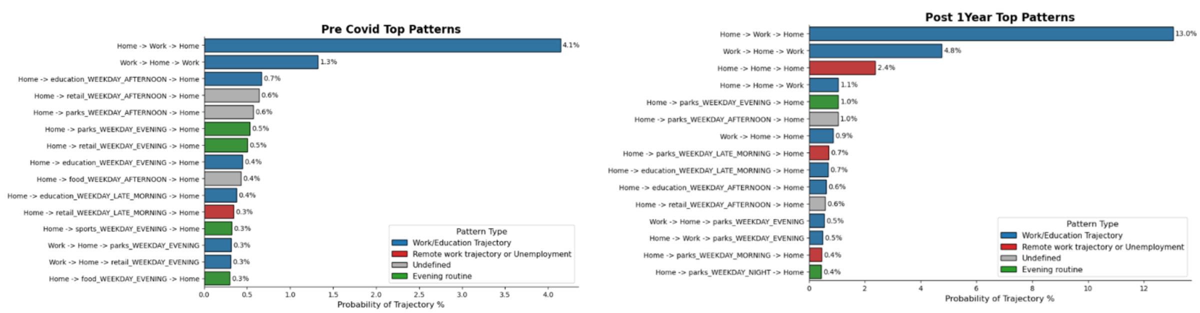

Figure 11: Top 3-grams extracted from the residents daily trajectories, comparing the 01/2019–01/2020 and 01/2022–09/2023 periods, colored by the type of the trip.

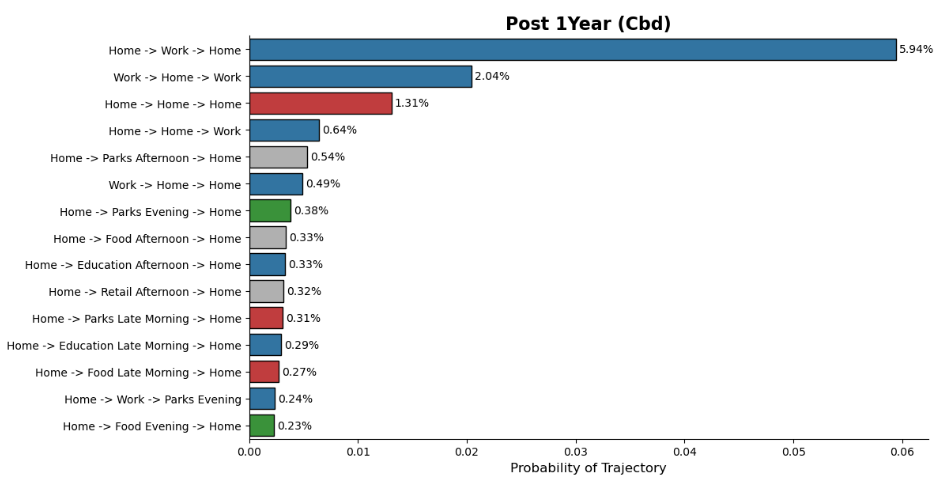

Figure 12: Top 3-grams extracted from the daily trajectories of residents whose detected home location falls within 5 km of the Azrieli Center, in 01/2022–09/2023 period, colored by the type of the trip.

The Word2Vec results show how close POI locations are to work locations in terms of their positions within daily trajectories before and two years after COVID-19. The identified list of places semantically closest to 'work' in 2019 includes transition points where people stop on the way to work, as well as public and religious places. As such, Donut Shops and Transition points have a cosine similarity of more than 50%, while Community Centers, Religious places, and Public Services have more than 40%. One unexpectedly high place is Entertainment spots, which also indicate a similarity higher than 40% (43.1%).

The list of places semantically close to 'work' in 2022 differs from the previous list. It includes educational places (59%), medical institutions (51%), and parks and sports locations (55%), as well as places suitable for remote work: Coffee Shops (55%) and Bakeries (54%). A t-SNE representation of the embeddings on aggregated tokens, built using `https://projector.tensorflow.org/`, is presented in Appendix C.

4.6 Stepwise Regression Results

Table 3 presents the results of the stepwise Two-Way Fixed Effects (TWFE) regressions across five distinct mobility and consumption outcomes (Panels A–E). By progressively introducing the distance gradient (M1), remote work potential (M2), local amenity dose (M3), and a saturated model with all interactions (M4), we isolate the drivers of post-pandemic behavioral shifts. Across the specifications, the coefficients for identified Third Places (Log_TP) reveal a significant, location-dependent effect on local retention, driven by consistent directional patterns in the interaction terms.

Panel A evaluates the share of local POI hours relative to all POI hours. Although the coefficients exhibit directional behavior similar to the subsequent models—specifically a positive main effect for Third Places and a negative spatial decay in the triple interaction—they do not reach statistical significance across the specifications (M1–M4). This suggests that we do not have enough observations to achieve statistical significance for the strict relative substitution between local and non-local commercial venues, even though the underlying behavioral pattern aligns with the broader trends observed in our other models.

In Panel B, we examine the allocation of local time toward commercial amenities. Focusing on the Third Place models, we find a significant positive coefficient for both distance ($Post \times Dist = 0.0496$) and Third Place availability ($Post \times Dose = 0.0482$ in M3). This indicates that a higher local density of Third Places correlates with residents spending more of their local time in neighborhood commercial spaces. However, the triple interaction ($Post \times Dist \times Dose$) yields a significant negative coefficient (-0.0161). This reveals a spatial decay: while Third Places have a localized attraction effect in the urban core, this effect

Table 3: Main Results: Stepwise WLS Regressions for Spatial & Commercial Retention (1 km Buffer)

Term	**M1 (Distance)**	**M2 (WFH_POT)**	**M3 (Amenity 1km)**		**M4 (Amenity&WFH_POT)**	
	Distance	Capacity	Log_Cand	Log_TP	Log_Cand	Log_TP
Panel A: Local POI hours / all POI hours						
Post × Dist	-0.0037 (0.0093)	-0.0198 (0.0308)	0.0291 (0.0461)	0.0185 (0.0259)	0.0218 (0.0482)	0.0090 (0.0296)
Post × WFH_POT		-0.0008 (0.0012)			-0.0012 (0.0052)	-0.0016 (0.0054)
Post × Dist × WFH_POT		0.0002 (0.0006)			0.0011 (0.0023)	0.0013 (0.0024)
Post × Dose			0.0397 (0.0293)	0.0343 (0.0265)	0.0429 (0.0385)	0.0393 (0.0361)
Post × Dist × Dose			-0.0060 (0.0123)	-0.0067 (0.0118)	-0.0069 (0.0155)	-0.0082 (0.0149)
Panel B: Local POI hours / all local hours						
Post × Dist	0.0065 (0.0046)	0.0148 (0.0140)	0.0804 (0.0219)***	0.0496 (0.0130)***	0.0912 (0.0233)***	0.0472 (0.0142)***
Post × WFH_POT		-0.0001 (0.0006)			-0.0025 (0.0027)	-0.0035 (0.0029)
Post × Dist × WFH_POT		-0.0002 (0.0002)			0.0005 (0.0011)	0.0010 (0.0012)
Post × Dose			0.0505 (0.0139)***	0.0482 (0.0128)***	0.0652 (0.0200)***	0.0660 (0.0195)***
Post × Dist × Dose			-0.0186 (0.0056)***	-0.0161 (0.0054)***	-0.0246 (0.0079)***	-0.0234 (0.0077)***
Panel C: Local candidate hours / all candidate hours						
Post × Dist	–	–	0.2046 (0.0622)***	0.1037 (0.0340)***	0.1589 (0.0637)**	0.0932 (0.0381)**
Post × WFH_POT		–			0.0080 (0.0067)	0.0058 (0.0069)
Post × Dist × WFH_POT		–			-0.0010 (0.0032)	-0.0002 (0.0033)
Post × Dose			0.1386 (0.0385)***	0.1305 (0.0344)***	0.0879 (0.0502)*	0.0960 (0.0464)**
Post × Dist × Dose			-0.0582 (0.0177)***	-0.0596 (0.0158)***	-0.0373 (0.0220)*	-0.0445 (0.0201)**
Panel D: POI hours / all hours						
Post × Dist	-0.0035 (0.0041)	0.0274 (0.0129)**	0.0675 (0.0151)***	0.0388 (0.0087)***	0.0817 (0.0151)***	0.0374 (0.0097)***
Post × WFH_POT		0.0008 (0.0005)			-0.0030 (0.0020)	-0.0040 (0.0021)*
Post × Dist × WFH_POT		-0.0006 (0.0002)***			0.0006 (0.0008)	0.0010 (0.0008)
Post × Dose			0.0501 (0.0097)***	0.0469 (0.0086)***	0.0683 (0.0134)***	0.0674 (0.0125)***
Post × Dist × Dose			-0.0170 (0.0039)***	-0.0160 (0.0035)***	-0.0244 (0.0053)***	-0.0245 (0.0050)***
Panel E: Radius of gyration, POIs (meters)						
Post × Dist	-453.4 (62.2)***	-620.8 (205.3)***	-838.3 (257.8)***	-708.2 (141.0)***	-635.2 (276.1)**	-618.3 (148.8)***
Post × WFH_POT		-12.62 (7.07)*			-23.3 (30.1)	-11.7 (31.6)
Post × Dist × WFH_POT		1.09 (3.37)			-2.1 (13.4)	-6.8 (14.1)
Post × Dose			-255.9 (152.3)*	-282.6 (132.6)**	-87.3 (240.7)	-189.0 (220.2)
Post × Dist × Dose			98.1 (68.5)	108.7 (60.5)*	26.3 (101.7)	63.4 (93.8)
Model Statistics						
Number of Observations	2,420	2,405	2,420	2,420	2,420	2,420
R-squared (Panel A)	0.5947	0.5952	0.5836	0.5830	0.5841	0.5834
R-squared (Panel B)	0.5866	0.5870	0.6993	0.7008	0.7008	0.7027
R-squared (Panel C)	–	–	0.5461	0.5471	0.5484	0.5491
R-squared (Panel D)	0.6410	0.6477	0.7776	0.7785	0.7804	0.7819
R-squared (Panel E)	0.7702	0.7752	0.7300	0.7304	0.7325	0.7327
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Statistical Area Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Evaluated at Dist = 0, where Dist is calculated as $\ln(Distance + 1)$.

*Format: Coef (SE). * p < 0.10, ** p < 0.05, *** p < 0.01. Panel C* $N = 2,402$.

diminishes as one moves further into the suburbs. These directional relationships and their significance remain robust in the fully saturated model (M4).

Panel C models the retention of time spent specifically in candidate work-substitute venues. The results echo the neighborhood vitality findings from Panel B. Near the urban core, Third Places help anchor residents to their immediate neighborhoods, indicated by the significant positive coefficient ($Post \times Dose = 0.1305$ in M3). The significant negative triple interaction (-0.0596 in M3) demonstrates that this localized retention effect fades along the distance gradient toward the metropolitan periphery. Both the main effects and the triple interaction remain statistically significant with the same directional signs in the saturated model (M4).

Panel D assesses the overall share of time dedicated to commercial venues regardless of location. The significant positive coefficient for Third Places (0.0469 in M3) indicates that these amenities stimulate overall consumption. Consistent with previous panels, the significant negative triple interaction (-0.0160 in M3) shows that the utility of these amenities is sensitive to the built environment, exhibiting stronger effects in the center compared to suburban contexts. These results are robust across both M3 and M4.

Finally, Panel E models the physical travel footprint (in meters) for commercial consumption. The baseline distance interaction ($Post \times Dist$) is negative and significant across all models (e.g., -708.2 meters in M3), showing that suburban residents shrank their travel radii post-pandemic compared to core residents. However, the negative coefficient for Third Places on reducing this travel radius ($Post \times Dose = -282.6$ meters) is only statistically significant in Model 3; it loses significance in the fully saturated model (M4) once occupational remote work capacity is controlled for. Furthermore, the triple interaction ($Post \times Dist \times Dose$) is not statistically significant in either M3 or M4, indicating that the spatial decay effect observed in the time-share models does not robustly apply to the overall physical radius of gyration.

5 Discussion of Results

The identification of post-pandemic third places reveals a distinct, yet spatially heterogeneous, shift in remote work behavior across the metropolitan area. While venues like coffee shops uniformly emerge as the most frequent work substitutes, their utilization patterns diverge significantly depending on the neighborhood context. In the urban core, third places exhibit a pronounced decrease in direct arrivals from the workplace, coupled with highly localized morning visit peaks, suggesting they function as direct, immediate office substitutes for hybrid workers. Conversely, in satellite cities, the behavioral signature shifts toward extended early-afternoon stays, with residents frequently utilizing alternative spaces like local libraries and cultural centers. This confirms that while the remote work behavioral pattern is clearly present in the data, its manifestation is non-uniform, adapting to the varying density, demographic needs, and built-environment compositions of different geographic zones.

The stepwise Two-Way Fixed Effects regressions further untangle this spatial heterogeneity, demonstrating that the magnetic pull of third places is highly localized and subject to steep spatial decay. The introduction of amenity doses—specifically the local density of identified third places—significantly increases the explanatory power for localized commercial vitality and overall out-of-home consumption. The models reveal a strong positive main effect, indicating that a robust local stock of third places effectively anchors residents, pulling time previously spent at home into the immediate neighborhood. However, the consistent, significant negative triple interaction between the post-pandemic shock, distance to the CBD, and amenity dose exposes a critical geographic limitation. It reveals that the localized utility of these amenities thrives in dense urban centers but struggles to replicate the same neighborhood retention power in suburban peripheries.

Evaluating the physical travel footprint through the radius of gyration solidifies this narrative of structural reorganization. Methodologically, analyzing travel footprint contractions via continuous distance interactions builds upon spatial decay frameworks commonly detailed in literature. In particular, spatial interaction gravity models are frequently applied in these studies to demonstrate how physical distance exponentially decays localized amenity attraction and reshapes urban equilibrium. Applying this framework to our regression outcomes, the baseline results show a massive, universal contraction in travel radii for suburban residents post-pandemic. Furthermore, a high dose of central third places acts as an independent anchor, significantly shrinking the travel radius by over 280 meters in the core. Yet, outside the urban core, this radius-shrinking utility is effectively neutralized. This ultimately suggests that while remote work has decentralized populations, suburban infrastructure remains inherently auto-dependent, forcing residents to maintain larger physical consumption footprints even when local third-place alternatives exist.

References

Barrero, J. M., Bloom, N., and Davis, S. (2023). The evolution of work from home. *Journal of Economic Perspectives*, 37:23–49.

Crivellari, A. and Resch, B. (2022). Investigating functional consistency of mobility-related urban zones via motion-driven embedding vectors and local poi-type distributions. *Computational Urban Science*, 2(1):19.

Esuli, A., Petry, L. M., Renso, C., and Bogorny, V. (2018). Traj2user: exploiting embeddings for computing similarity of users mobile behavior.

Gokan, T., Kichko, S., Matheson, J., and Thisse, J.-F. (2022). How the rise of teleworking will reshape labor markets and cities. *SSRN Electronic Journal*.

Li, D., Zheng, Y., Wang, S., Guo, X., and Zhao, J. (2026). Quantifying the nonlinear causal impact of commute time on u.s. remote work. *Transportation Research Part D: Transport and Environment*, 151:105153.

Li, W., Zhang, E., and Long, Y. (2024). Unveiling fine-scale urban third places for remote work using mobile phone big data. *Sustainable Cities and Society*, 103:105258.

Mariotti, I., Capdevila, I., and Lange, B. (2023). Flexible geographies of new working spaces. *European Planning Studies*, 31(3):433–444.

Murray, D., Yoon, J., Kojaku, S., Costas, R., Jung, W.-S., Milojević, S., and Ahn, Y.-Y. (2023). Unsupervised embedding of trajectories captures the latent structure of scientific migration. *Proceedings of the National Academy of Sciences*, 120(52).

Oldenburg, R. (1989). *The Great Good Place: Cafés, Coffee Shops, Bookstores, Bars, Hair Salons, and Other Hangouts at the Heart of a Community*. Marlowe & Company, 2nd edition.

Tomaz, E. and Tabrizi, H. A. (2024). The evolution of non-traditional workplaces: From third places to hybrid places. In *SpringerBriefs in Applied Sciences and Technology*, pages 7–20.

Tregubova, I. and Gdaliahu, M. (2025). Estimating the impact of working from home on urban equilibrium: neighborhood scale effects using mobile data.

Weng, G., Kim, M., Ahn, Y.-Y., and Moro, E. (2025). Beyond distance: Mobility neural embeddings reveal visible and invisible barriers in urban space. *arXiv*.

Yang, M., Zheng, Y., Wang, S., Guo, X., and Zhao, J. (2025). Impacts of the built environment on remote work choices and preferences. *Cities*, 163:106034.

Zontag, N., Madhala, S., and Bental, B. (2022). Working from home in israel. Technical report, Taub Center for Social Policy Studies in Israel.

A Colleced POIs Categories

POI categories	**POI keywords**
Health	hospital, clinic, pharmacy, dentist, doctors
Education	school, university, college, kindergarten
Entertainment	cinema, theatre, nightclub, arts_centre, community_centre, library, museum
Food	bar, cafe, restaurant, fast_food, pub, food_court
Outdoor attractions	attraction, memorial, ruins, artwork, fountain, archaeological
Parks and sport	park, pitch, sports_centre, stadium, playground, swimming_pool, dog_park, viewpoint
Retail	supermarket, mall, convenience, clothes, pharmacy, bakery, electronics, furniture, books, shoes, jewelry, sports, toys, optician, stationery, department_store, florist, kiosk, car_dealership
Services	bank, hairdresser, nails, spa, massage, tattoo, photography, travel_agency, bicycle_repair, car_repair, car_wash, post_office, bicycle_rental

Table 4: POI Categories and Associated Keywords

B Keywords extracted from Google Reviews

Exact direct remote-work keyword list used in the review matching:

- "laptop",
- "computer",
- "business meeting",
- " place for work",
- " place to work",
- " get the work done",
- "remote work"

And the exclusion list (reviews containing these are filtered out):

- events
- funeral
- repair
- wifi doesn't work
- on the go
- doctor
- wifi only
- no wifi

Keywords describing atmosphere suitable for remote working used to create a metric flag_atmosphere_keyword:

- " peace"
- " quiet"
- " relaxing"

- ”relax”
- ” comfortable”
- ” calm”
- ”meeting friends”
- ”sit with”
- ”sitting with”
- ”wifi”
- ”wi-fi”
- ”charger”
- ”socket”

And the exclusion list:

- ”not quiet”
- ”rush”
- ”loud”
- ”not comfortable”
- ”not relaxing”
- ”not calm”

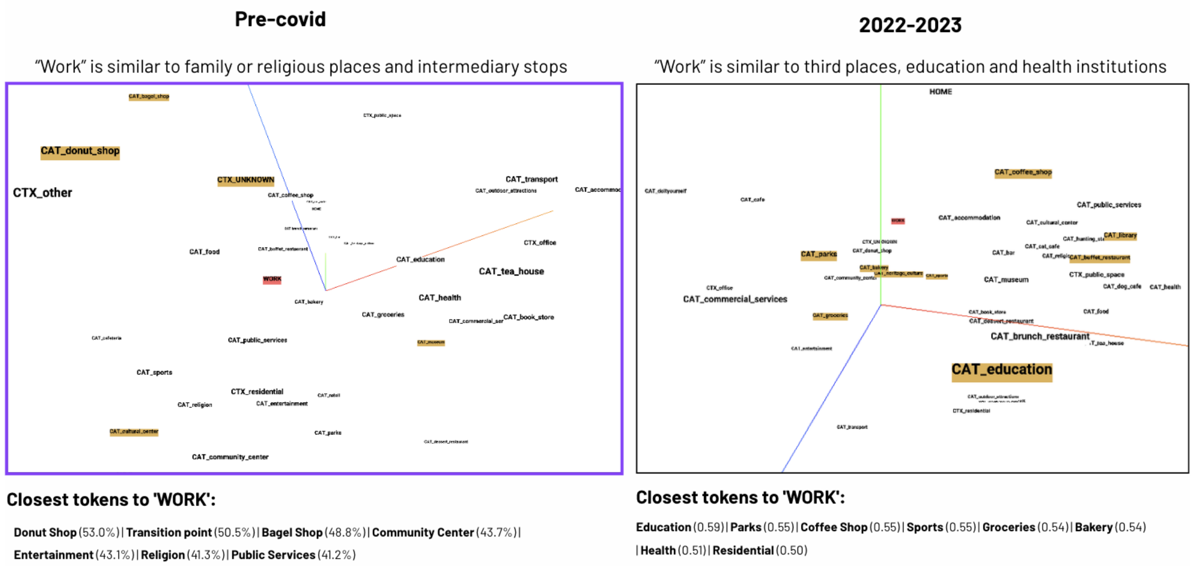

Figure 13: Top 3-grams extracted from the daily trajectories of residents whose detected home location falls within 5 km of the Azrieli Center, in 01/2022–09/2023 period, colored by the type of the trip.

C T-SNE representation of Word2Vec embeddings

T-SNE visualization of the Word2Vec embeddings build on pre-treatment sample. Locations that are semantically related to the 'work' are highlighted within the vector space