The Politics of Visibility: Party Symbols, Alliances, and Electoral Returns in Istanbul's *Mahalle*s*

Cantay Caliskan[1]**

University of Rochester, Rochester, NY, USA

Abstract. What citizens see on their morning commute—rows of party flags strung between buildings, banners draped across storefronts—turns out to carry measurable electoral weight. Using 748,783 Google Street View panoramas from every precinct (*mahalle*) in Istanbul, we identified roughly 1.22 million partisan symbols and compared their spatial distribution with official election results. The pattern is striking: when a district's streets host an additional 10,000 flags or banners, the sponsoring party's vote share increases by about two to four percentage points, while overall turnout remains effectively unchanged. Governing parties play defense, crowding their banners into an average footprint of just $0.027\,\mathrm{km}^2$ per flag—AKP symbols cram into an extreme $0.007\,\mathrm{km}^2$—while the opposition stretch each flag over roughly $0.10\,\mathrm{km}^2$, and the smallest parties span whole tens of square kilometres. Formal coalitions leave no discernible street-level trace—the spacing between partners' symbols is no tighter than between rivals—and neighborhood traits such as income, building density, or land use remain poor predictors of where the fabric ultimately hangs. The sole exception is age: precincts with a larger share of residents in their early twenties display noticeably denser layers of visual propaganda. By transforming millions of ordinary street images into a city-wide map of political visibility, this study offers election officials, journalists, and scholars a novel, scalable tool for tracking how the battle for votes unfolds in public space.

Keywords: Visual Politics · Political Symbols · Turkish Politics · Computational Social Science

"(People) generally judge more by the eye than by the hand, for everyone can see and few can feel."

— Niccolò Machiavelli, *The Prince*

* The Turkish term *mahalle* derives from Arabic *maḥalla(t)*, "lodging place" or "abode," following the *mafala(t)* pattern (noun of time and place) from the verb *ḥalla*, "to settle." For etymological details, see the entry *maḥalle*.

** I thank Anna Fang, Aryan Raj Dhawan, Brendan McCollum, Danielle Colelli, and Xiaoyu Zhao for outstanding research assistance. I also thank Prof. Mert Moral from Sabancı University for providing the extensive data on election results from Turkey.

1 Introduction

At sunrise in Istanbul, ferries thread the Bosphorus while ribbons of colourful party flags—orange-white bulbs of the *Adalet ve Kalkınma Partisi* (AKP), six red arrows of the *Cumhuriyet Halk Partisi* (CHP), crescent trios of the *Milliyetçi Hareket Partisi* (MHP)—flutter above still-shuttered shops. During election months these pennants multiply overnight, turning ordinary streets into kilometre-long plebiscites in which commuters cast their first "vote" with their eyes. The campaign spectacle soon escalates: sound trucks pump anthems through alleyways, cranes drape leader portraits across arterial roads, and rival crews wage a nightly contest of removal and re-hanging. The city's built environment thus becomes a contested canvas on which narratives of nation, secularism, faith, and class are painted banner by banner—a visual politics that has come to define Turkey's politics since its foundation.

This paper enters that story at street level. We assemble nearly three-quarters of a million Google Street View panoramas spanning every election month since 2000 and apply state-of-the-art computer-vision models to detect more than 1.2 million partisan symbols. Linking these geolocated flags and banners to precinct-level returns, we ask how visibility in the urban landscape maps onto electoral fortunes and what this cartography of symbols reveals about the strategic logic of Turkey's governing and opposition blocs.

Since the early 2000s, polarization has mapped onto two broad coalitional camps that crystallized after the 16 April 2017 constitutional referendum, which narrowly approved a shift to an executive presidency (51.4% "Yes") and accelerated the concentration of power in the presidency.(OSCE/ODIHR, 2017; Öktem and Akkoyunlu, 2016; Esen and Gümüşcü, 2016; Aydın-Düzgit and Balta, 2017) Subsequent cycles combined very high turnouts by comparative standards with mounting doubts about procedural fairness and media pluralism, deepening the perception of a competitive–authoritarian equilibrium.(Kalaycıoğlu, 2022; Esen and Gümüşcü, 2016; Aydın-Düzgit and Balta, 2017) The referendum result (51.4–48.6, on turnout above 80%) inaugurated the institutional framework under which later alliances competed.(OSCE/ODIHR, 2017)

The governing *People's Alliance* (*Cumhur İttifakı*) formed ahead of the 2018 contests around the Justice and Development Party (AKP) and the Nationalist Movement Party (MHP), marrying conservative–Islamist incumbency to hard-nationalist rhetoric and yielding a durable governing bloc that carried through 2018 and 2023.(Denizeau, 2023) The opposition responded with the heterogeneous *Nation Alliance* (*Millet İttifakı*), led by the Republican People's Party (CHP) and the Good Party (İYİ) with smaller partners, unified by a pledge to restore parliamentary checks and curb presidential overreach.(Denizeau, 2023) The pro-Kurdish HDP—later fielding candidates under the Green Left banner—remained outside both blocs under sustained legal pressure, yet frequently acted as a pivotal extra-bloc force in close races.(Aydın-Düzgit and Balta, 2017; Esen and Gümüşcü, 2016) These confrontations produced watershed local outcomes: the 31 March 2019 Istanbul mayoral election was annulled on 6 May and re-run on 23 June, when CHP's Ekrem İmamoğlu prevailed decisively, under-

scoring the stakes and volatility of urban competition.(Reuters, 2019; Wuthrich and Ingleby, 2020)

These coalitional lines overlay long-standing ideological clusters. The AKP articulates a national-conservative, majoritarian project centered on strong leadership and developmental populism; the MHP anchors uncompromising Turkish nationalism; the CHP blends secular–Kemalist state tradition with a broadened social-democratic appeal; İYİ represents a secular, pro-NATO nationalist current that split from the MHP; and the HDP advances a left-pluralist platform emphasizing minority rights, gender equality, and decentralization.(Denizeau, 2023; Aydın-Düzgit and Balta, 2017; Esen and Gümüşcü, 2016) Personalization around leaders, securitized frames, clientelistic mobilization, and media asymmetries have tightened in-group identities and raised the costs of cross-camp persuasion.(Aydın-Düzgit and Balta, 2017; T. S. Baykan and Somer, 2022; Esen and Gümüşcü, 2016)

In this context, the politics of visibility has grown central. Campaigns routinely transform streetscapes into partisan canvases—strings of party pennants, leader portraits on facades, large banners at intersections, and ubiquitous national flags—while online echo chambers reinforce affective divides.(Kutlu et al., 2019) At the macro-symbolic scale, landmark interventions have inscribed ideological claims into Istanbul's built environment: most prominently, the reconversion of Hagia Sophia to a mosque on 10 July 2020 and the inauguration of Taksim Mosque on 28 May 2021 opposite the Republic Monument—widely read as asserting Islamist–national symbolism over historically secular civic space.(Sofos, 2021; Press, 2021; Tokdoğan, 2024a; Walton, 2024) These gestures coexist with bottom-up visual repertoires: the 2013 Gezi Park cycle crafted an inclusive, party-flag-free civic aesthetic to counter partisan saturation.(Aydın-Düzgit and Balta, 2017; Walton, 2024)

Our study situates Istanbul's *mahalle* streets within this national tableau by quantifying visual saturation and linking it to ballot returns. From 748,783 Google Street View panoramas taken during election and pre-election months since 2000, we extract about 1.22 million partisan symbols. We find that an extra 10,000 flags or banners in a district raises the sponsor's vote share by roughly two to four percentage points, yet has no systematic effect on turnout, indicating that visibility persuades more than it mobilizes. Revisiting spatial strategy, we show that clustering is driven mainly by organizational capacity. Well-financed parties—both incumbents (AKP, MHP) and the principal opposition (CHP)—saturate their strongholds, averaging one flag per 0.01 km^2 (AKP: 0.007; CHP: 0.010). Across blocs, government parties occupy 0.027 km^2 per symbol, while resource-constrained challengers stretch each banner over about 0.10 km^2, and fringe groups over double-digit square kilometers. Finally, parties disproportionately target youth-heavy precincts: the density of symbols rises markedly where residents aged 20–25 form a larger share of the electorate. By converting millions of ordinary street images into a city-wide map of political visibility, the paper links Turkey's coalitional polarization to the everyday spatial

economy of symbols and offers a scalable template for tracking visible politics under democratic backsliding.

1.1 Emotional Politics in Turkey

Emotion lies at the heart of this process. Political symbols—flags, colors, leader portraits—bind adherents through shared affect, shaping collective memory and national identity. These mechanisms resonate far beyond Turkey. Pakistani parties weaponize ballot symbols to mobilize illiterate voters (Ahmad, 2020); India's iconic party emblems sway vote shares even among independents (Lee and Qi, 2024); removal of Soviet-era insignia in Ukraine boosted support for parties nostalgic for the USSR (Rozenas and Vlasenko, 2022); and Ukraine's *Orange Revolution* famously coalesced around a single color (Wilson, 2005). Such cases underscore Edelman's enduring claim that politics is "a contest of symbols" (Edelman, 1964). In fragmented societies, meanings are disputed and visual space becomes a literal battleground over power and belonging.

Political actors in Turkey have long recognized that visibility in public space translates into political influence. In urban environments such as Istanbul, symbols, flags, and banners serve not only as campaign advertisements but as instruments of spatial communication. Scholars note that public space is "the strongest way for social groups to show themselves to larger audiences and to be recognized as part of society" (Mitchell, 1995).

During election periods, Turkish parties transform cityscapes into symbolic battlegrounds. Streets become adorned with AKP's orange–white pennants, CHP's red flags with six arrows, and the İYİ Party's yellow sunburst emblems. These visual markers convey not just party identity but signal territorial claims, organizational capacity, and voter mobilization. In contested districts, the volume and placement of party flags indicate dominance and often lead to "wars of symbols," in which rival groups vie for visual superiority (T. S. Baykan and Somer, 2022).

Beyond flags, monuments and naming practices form part of the state's symbolic apparatus. Early-Republican reforms sought to "de-Ottomanize" the urban landscape, erecting statues, plazas, and institutions that asserted secular-nationalist ideology (Atakuman, 2010). Taksim Square—anchored by the 1928 Republic Monument—became the new civic heart of Istanbul precisely through such spatial engineering (Aysegul Baykan and Hatuka, 2010).

The AKP has reversed this trajectory by embedding Ottoman-Islamist symbolism into the city's geography. The Taksim Mosque, inaugurated in 2021 opposite the Republic Monument, exemplifies a "spatial victory" over a historically secular site (Tokdoğan, 2024b). Parallel gestures include renaming the Bosphorus Bridge as the "15 July Martyrs' Bridge" after the failed 2016 coup and, most spectacularly, reconverting Hagia Sophia into a mosque in 2020 (Sofos, 2021).

Mass gatherings likewise appropriate public space. Post-coup "democracy watches" filled squares with national flags and portraits of Erdoğan, staging regime support as national unity (Karaveli, 2017). By contrast, the 2013 Gezi

Park protests rejected party insignia, instead adopting civic and inclusive symbols that turned a small green patch into a nationwide emblem of dissent (Tuğal, 2013).

Finally, megaprojects—from the Çamlıca Mosque to Istanbul Airport and the network of "Nation's Gardens" (*millet bahçeleri*)—function as monumental assertions of ideological dominance and technological prowess, buttressing the government's populist narrative of national rebirth (Erensü et al., 2022).

1.2 Street-Level Imagery for Social and Political Behavior

Timnit Gebru and collaborators were among the first to demonstrate the social-scientific value of Google Street View (GSV) imagery. In a landmark study, they trained a fine-grained deep-learning pipeline that detected 22 million cars in 50 million GSV photos from 200 U.S. cities and showed the resulting "virtual vehicular census" could recover census-tract income, racial composition, education, and presidential vote share with high accuracy (Gebru et al., 2017). Their work advanced the idea that routinely collected street photographs can serve as a near-real-time visual census capable of supplementing traditional household surveys.

Building on this insight, researchers have used GSV to study an expanding set of social and political outcomes. Freitas *et al.* combined deep computer vision with crowdsourced ratings to quantify perceived safety and maintenance across Raleigh, North Carolina, showing that shifts in urban perception track post-recession socioeconomic changes (Freitas et al., 2023). Cinnamon and Jahiu's global review catalogues more than 600 studies that leverage panoramic street imagery for topics ranging from public health to political geography, while also noting limits in temporal coverage and corporate data control (Cinnamon and Jahiu, 2022).

A complementary line of inquiry models how streets look to human observers. Naik *et al.* introduced *StreetScore*, predicting perceived safety for one million streetscapes crowd-rated via the Place Pulse platform (Naik, Philipoom, et al., 2014). Their follow-up, *StreetChange*, compared historical panoramas of the same block to measure visual improvement between 2007 and 2014, finding that college-educated populations, higher baseline visual quality, and downtown proximity predict gentrification-like upgrades (Naik, Kominers, et al., 2017).

Outside the United States, Suel *et al.* trained end-to-end convolutional networks on 500,000 London images to estimate income, education, and living environment, then demonstrated that the model generalizes—after minimal retraining—to other U.K. cities (Suel, Barros, et al., 2019). In a twelve-city comparison across five countries, the same authors showed that wealthy neighborhoods share a common visual signature, whereas disadvantaged areas remain city-specific, cautioning against naïve cross-city transfer of poverty models (Suel, Kolding, et al., 2023). At the micro-scale, Ilic *et al.* detected façade renovations across multi-year GSV sequences in Ottawa to map block-by-block gentrification patterns (Ilic et al., 2019).

Collectively, these studies illustrate how street imagery paired with modern machine learning offers scalable, low-cost proxies for neighborhood demographics, political affiliation, inequality, and change. As GSV coverage expands and computer vision advances, the technique promises increasingly granular and global insights into the social and political life of cities.

1.3 Research Contribution

This article presents the first large-scale, street-level visual analysis of how visible partisan symbolism correlates with electoral outcomes in Turkey. We compile **747,783** Google Street View panoramas captured in Istanbul during and immediately before local and general election months. Leveraging state-of-the-art computer-vision models, we detect and classify party flags, banners, and logos, thereby producing a high-resolution "symbolic map" of the city. Merged with precinct-level vote returns, these data allow us to test whether variation in symbol density predicts ballot-box performance—an approach that extends earlier work using Street View to infer socioeconomic patterns (Gebru et al., 2017) into the realm of political behavior.

Substantively, our findings illuminate how spatial propaganda operates in an environment defined by permissive outdoor-display regulations, which allow campaigns to saturate public streetscapes with party symbols (for an example, see Figure 1). Methodologically, the project demonstrates the value of combining computer vision with political science theory to study affective polarization and symbolic politics at scale. More broadly, the framework is exportable to other settings where democratic backsliding and partisan spectacle intersect—from Hungary's poster wars to Brazil's color-coded rallies. By quantifying the battle of symbols in the streets, we shed new light on the visual terrain of contemporary electoral competition.



Fig. 1: Political flags on streets of Istanbul (example).

2 The Architecture of Power in Turkish Political Parties: Centralization and Leadership Dominance

Turkish parties are, by historical design and contemporary practice, highly centralized, top-down organizations. Scholarship consistently describes "oligarchic" structures in which leaders and central executive boards dominate nominations, promotions, programmatic choices, and—crucially for campaigns—the allocation of financial and material resources.(P. Ayan, 2010; T. S. Baykan and Somer, 2022) This logic spans the spectrum: the AKP is frequently characterized by hegemonic intra-party control, while the CHP exhibits a more oligarchic form; in both, decisive authority resides at the *Genel Merkez*. The closed-list proportional representation system further entrenches top-down discipline by tying candidates' careers to centrally determined list placement rather than to geographically localized constituencies.(P. Ayan, 2010; Özbudun, 2000)

Centralization is also formally codified. The AKP statute authorizes the Central Executive Committee to take "all necessary decisions regarding all elections," to set "working methods and strategies," and to administer election work—textual authority that mirrors academic diagnoses of leadership dominance.(Adalet ve Kalkınma Partisi (AK Party), 2015) Historically, the CHP's origins as a vanguard party of state-led modernization normalized centralized command, a pattern reproduced by later parties including the DP and its successors.(Somer, 2015; Heper, 1998; Plaggenborg, 2015; P. Ayan, 2010)

2.1 The Mechanics of Campaign Management and Resource Allocation

Modern campaigns are professionalized operations led from party headquarters by coordinators, strategists, pollsters, and communications teams who define national messaging and budgets; provincial (*il*) and district (*ilçe*) branches execute rather than design strategy, ensuring message discipline country-wide.(Göksu, 2019; Adalet ve Kalkınma Partisi (AK Party), 2015) Clientelism provides the key linkage on the ground. Turkey's system has evolved from notable-based patronage into institutionalized, party-broker clientelism. The AKP in particular has been analyzed as a "national machine" for partisan distribution of public resources, with local activists identifying needs, relaying them up the hierarchy, and delivering benefits in exchange for support.(Sayarı, 2014; Esen and Gümüşcü, 2016; Somer, 2019; Buğra and Candas, 2011; T. S. Baykan and Somer, 2022)

Symbols are part of this centrally managed allocation. Optimization models developed for the Turkish context explicitly treat "party flags" as an advertising line item to be distributed alongside television and social-media spending to maximize votes across places—evidence that street-level visibility is planned, budgeted, and routed, not merely spontaneous enthusiasm.(Akgün, 2016)

This hierarchy also sustains a form of "soft" or symbolic clientelism: dense, highly visible flagging in a neighborhood publicly signals organizational capacity

Table 1: The hierarchical tiers of Turkish election campaign management

Level	Key Actors	Primary Role
Party Headquarters (*Genel Merkez*)	Party Leader, Central Executive Committee, Campaign Coordinator	Strategy formulation; national messaging; macro-resource allocation; candidate approval; media strategy
Provincial Organization (*İl Teşkilatı*)	Provincial Chair, Executive Board	Regional coordination; distribution of centrally-allocated resources; large-scale rally organization; media relations
District Organization (*İlçe Teşkilatı*)	District Chair, Executive Board	Local execution; distribution of materials (flags, posters); door-to-door canvassing; precinct-level monitoring
Volunteers & Clientelistic Networks	Party members, local notables, beneficiaries of patronage	Ground-level visibility; displaying symbols; mobilizing voters; demonstrating local support

and territorial authority—effectively, "we provide here"—thereby complementing material inducements and, at the margin, deterring rivals.(Tokdoğan, 2024b)

2.2 Istanbul as a Symbolic Battlefield: The Macro-Politics of Urban Space

Parties approach competition through a political-marketing lens in which voters are consumers and leaders/platforms are brands. Within this framework, symbols, music, and colors are core tools to build institutional image, secure visibility, and persuade undecided voters.(Cwalina et al., 2011; Lees-Marshment, 2014; Niffenegger, 1989) Short-term campaign displays interact with a long-term state project of inscribing ideology onto the city. Since the 2000s, conservative-Islamic and neo-Ottoman symbolism has been advanced through monumental interventions—renaming the Bosphorus Bridge as the "15 July Martyrs' Bridge," reconverting Hagia Sophia to a mosque (10 July 2020), and inaugurating Taksim Mosque (28 May 2021) opposite the Republic Monument—casting Istanbul's built environment as a battleground for politics.(Tokdoğan, 2024a; Walton, 2024; Ayşegül Baykan and Hatuka, 2010)

Centralized command can also misread local sentiment. Analysts link the AKP's municipal setbacks in 2019 and its broader local defeat in 2024 to organizational rigidity and the opposition's use of locally resonant candidates—İmamoğlu in Istanbul and Yavaş in Ankara—who won by 54–45 in the June 23, 2019 rerun and again by clear margins on March 31, 2024.(Reuters, 2019; Wuthrich and Ingleby, 2020; Balta, 2024)

2.3 Justifying a Quantitative Analysis of Political Symbolism

Three pillars justify treating flag density as a non-random independent variable. First, Turkish parties are highly centralized "oligarchic" organizations (B. Ayan, 2010; Somer, 2021). Second, campaigns are executed through hierarchical logistics in which headquarters devise strategy and local branches carry it out, often via clientelistic networks (Göksu, 2019; Akgun, 2016; Somer, 2019).

Third, micro-level flag placement aligns with a macro-level strategy of spatial symbolism (Tokdoğan, 2024b; Dorroll, 2024).

Because the decision to place symbols is deliberate and data-informed, the resulting spatial pattern is a tangible signal of organization and intent. Far from "noise," each flag constitutes political data—data suitable for rigorous statistical testing of hypotheses such as the link between symbol density and vote share.

3 Major Examples of Visual Symbolism in Contemporary Turkish Politics

Political communication in Turkey is mediated as much through imagery as through rhetoric. Urban streetscapes are densely populated with symbols that encode ideological commitments, historical narratives, and claims to civic space. The discussion below surveys the best-known emblems of Turkey's five largest political parties, followed by the supra-partisan symbolism of Mustafa Kemal Atatürk and the national flag. Throughout, we highlight design features that later sections will operationalize for automated detection.

Justice and Development Party (AKP) The Justice and Development Party (*Adalet ve Kalkınma Partisi*, AKP) relies on a three-color palette—white, orange, and blue—with orange providing the primary chromatic emphasis. Its core logotype, a tungsten-yellow light bulb on a white ground, circulates on building signage, street garlands, and digital media. The illuminated bulb evokes innovation and forward momentum while subtly echoing Qur'anic metaphors of divine light, thereby linking the party's technocratic modernism to its Islamic lineage.

Republican People's Party (CHP) The Republican People's Party (*Cumhuriyet Halk Partisi*, CHP) deploys a single canonical device: six white arrows radiating from a common point on a red field. The emblem appears in horizontal and vertical orientations and, less frequently, in color inversions. The Six Arrows represent the foundational principles of Kemalism—Republicanism, Populism, Nationalism, Laicism, Statism, and Reformism—thereby inscribing the party's ideological heritage directly onto its visual identity.

Peoples' Democratic Party (HDP) The Peoples' Democratic Party (*Halkların Demokratik Partisi*, HDP) employs a stylized tree with multicolored leaves, set against a purple-and-green ground. Two minor variants differ only in whether the party name appears below the icon. The arboreal motif foregrounds the HDP's pluralistic coalition—Kurdish, leftist, and environmentalist constituencies— and invokes mythic references to the "Tree of Life," underscoring themes of continuity and collective growth.

Nationalist Action Party (MHP) The Nationalist Action Party (*Milliyetçi Hareket Partisi*, MHP) exhibits strict visual discipline, limiting its flag repertoire to ruby red and white. Its primary emblem—three white crescents, occasionally accompanied by the MHP monogram—adapts an Ottoman martial standard to convey strength, pan-Turkic unity, and decisive leadership. The design has remained virtually unchanged since 1969, reinforcing the party's emphasis on continuity and tradition.

Good Party (İYİ) The Good Party (*İYİ Partisi*) employs a yellow sun and the white "İYİ" logotype on a cyan field, a motif derived from the Kayı tribal tamga. The sun connotes resilience, clarity, and renewal, while the bright blue-yellow-white triad positions the party as an optimistic, centrist force anchored in Turkic heritage yet oriented toward reform.

Atatürk and the National Flag Beyond partisan iconography, two supra-political symbols dominate Turkey's visual landscape: the national flag—a red field bearing a white crescent and star—and the ubiquitous portrait of Mustafa Kemal Atatürk. Frequently, the two are combined in a single banner that superimposes Atatürk's forward-gazing likeness on the flag's crimson ground. This composite synthesizes the founding ethos of the Republic with an emblem of sovereign unity, making it a powerful—and seemingly nonpartisan—resource for political actors across the ideological spectrum. Secular groups emphasize the image to underscore Kemalist values, whereas conservative parties employ it to affirm their attachment to the Republican legacy. Its high contrast and consistent color scheme render the montage a salient and easily detectable class in computer-vision applications.

4 Research Questions and Hypotheses

Building on visual–political scholarship that treats street-level iconography as an active campaign instrument, this study leverages Google Street View imagery, precinct variables, and district results to investigate how party symbols in Istanbul both reflect existing political support and actively shape electoral competition. To structure this inquiry, we pose four interrelated research questions, each linked to a core hypothesis.

RQ1: Visual Presence and Electoral Performance. To what extent is the *density* of a party's symbols in a precinct associated with that party's vote share and overall voter turnout?

RQ2: Symbolic Diversity and Electoral Competition. Do neighborhoods that exhibit a richer mix of multiple parties' symbols experience more competitive electoral races (e.g., smaller victory margins)?



(a) Atatürk + Flag



(b) AKP



(c) CHP



(d) HDP



(e) İYİ



(f) MHP

Fig. 2: Core visual symbols in Turkish politics: Atatürk iconography and major-party flags.

RQ3: Strategic Spatial Deployment. What do the specific spatial patterns of symbol placement—such as co-location with allies, clustering in strongholds, or dispersion across wider areas—reveal about different parties' campaign strategies?

RQ4: Socio-Demographic Targeting. How do parties strategically target their visual campaigns in relation to the socio-economic characteristics and built environment of different neighborhoods?

To guide the empirical analysis, we advance four corresponding hypotheses, each with testable sub-hypotheses that draw on insights from political communication, symbolic politics, and spatial electoral studies.

H1: Visual presence is positively associated with electoral performance.

H1a: *Across parties, greater symbol density in a district—normalized by the number of eligible voters and panorama coverage—is positively associated with that party's vote share.*
(Rationale: Normalizing for image-sampling intensity helps isolate the association between visual saturation and electoral support.)

H1b: *Across districts, greater overall symbol density is positively associated with higher voter turnout.*
(Rationale: Intense visual saturation from all parties may galvanize the electorate by making the election more salient, thereby boosting overall participation.)

H2: Symbolic diversity signals a competitive electoral environment.

H2a: *Precincts displaying a comparable share of symbols from three or more parties have narrower victory margins than precincts dominated by a single party's symbols.*
(Rationale: Multiple parties invest heavily in visibility where the outcome is most in doubt, generating a richer symbolic tapestry and, consequently, tighter races.)

H3: Spatial patterns of symbol placement reveal distinct campaign strategies.

H3a: *Allied parties' symbols (e.g., CHP and İYİ) co-locate more frequently and at shorter inter-symbol distances than non-allied pairs, net of overall density.*
(Rationale: Coalition partners often coordinate campaign activities and share supporter infrastructure, resulting in a measurable spatial clustering of their visual propaganda.)

H3b: *Dominant incumbent parties use a "defensive clustering" strategy by concentrating symbols in their strongholds, while challenger parties use an "offensive dispersion" strategy to maximize visibility and create an impression of broad support.*
(Rationale: An established party signals territorial control by saturating its core areas, whereas a challenger must spread its resources more thinly to build name recognition and reach swing voters.)

H4: Symbol placement is targeted based on local demographics and environment.

H4a: *Precinct-level flag counts for major parties covary systematically with local socio-economic structure—specifically with the age profile, educational attainment, housing stock, and density of commercial activity.* (Rationale: Parties are expected to concentrate visual symbols where demographic and economic conditions promise the highest marginal return on visibility, such as in dense residential areas or commercial corridors, adapting placement to the composition of the local electorate.)

5 Data

5.1 Training Data Collection for Political Symbols

Our guiding assumption was that party flags and banners are displayed most densely in the immediate vicinity of party offices; concentrating on those locations would therefore maximize the yield of clear, well-framed symbols for training. To operationalize this strategy, we first assembled a nationwide inventory of party offices. Using the Google Search API, we queried every combination of all available political parties (47 in total) across Turkey's 960 districts (provincial centers included). The initial crawl returned 218,895 address records—an unrefined set containing duplicates and misclassifications. We retained only those entries whose address label contained both the party's full name (or acronym) and the district name, yielding 14,583 verified office locations with precise latitude–longitude coordinates.

For each verified office we generated twelve probe points on a circle of radius 20 m: six equally spaced bearings (every 60°) at distances of 10 m and 20 m. Each probe was queried against Google Street View; when imagery existed, we stored the corresponding metadata. This procedure identified 379.158 panoramas in the immediate vicinity of party offices and—within the election and pre-election months—an additional 6.496 panoramas that captured heightened campaign activity. All 18.035 visible party symbols observed in those panoramas were manually annotated with the *Roboflow* platform, producing the largest annotated corpus of Turkish party symbols collected to date. Inter-annotator agreement (Fleiss' κ) exceeded 0.82. A detailed breakdown of the captured objects appears in Table 2.

5.2 Urban Landscape Panorama Collection for Inference

We relied on the Google Maps API to assemble the panoramic imagery used for training and evaluation. The API mirrors the continuously refreshed data available through the public interface, incorporating satellite updates, Street View vehicles, and crowd contributions that add more than twenty million data points each day (Russell, 2019). To delimit the universe of interest, we obtained the shapefile for the official polygon of Istanbul Province from the Turkish General

Table 2: Annotated Label Counts in the Training Dataset

LABEL	COUNT	TURKISH EXPLANATION	ENGLISH TRANSLATION
Turkish Flag	4272	Türkiye Bayrağı	Turkish Flag
AKP	3161	Adalet ve Kalkınma Partisi	Justice and Development Party
CHP	3135	Cumhuriyet Halk Partisi	Republican People's Party
SAADET	2289	Saadet Partisi	Felicity Party
MHP	1888	Milliyetçi Hareket Partisi	Nationalist Movement Party
IYI	822	İYİ Parti	Good Party
HDP (1)	410	Halkların Demokratik Partisi	Peoples' Democratic Party
BBP	396	Büyük Birlik Partisi	Great Unity Party
IP	292	İşçi Partisi	Workers' Party
YRP	215	Yeniden Refah Partisi	New Welfare Party
VP	205	Vatan Partisi	Patriotic Party
ATATURK	116	Atatürk	Atatürk (Founder of Modern Turkey)
DP	106	Demokrat Parti	Democratic Party
DSP	93	Demokratik Sol Parti	Democratic Left Party
BTP	92	Bağımsız Türkiye Partisi	Independent Turkey Party
DEVA	74	Demokrasi ve Atılım Partisi	Democracy and Progress Party
MP (1)	70	Memleket Partisi	Homeland Party
TKP	69	Türkiye Komünist Partisi	Communist Party of Turkey
YSP	57	Yeşil Sol Parti	Green Left Party
ZP	36	Zafer Partisi	Victory Party
UO	30	Ülkü Ocakları	Idealist Hearths
HYP	29	Halkın Yükselişi Partisi	People's Ascension Party
GW	22	Grey Wolves	Grey Wolves
MP (2)	18	Millet Partisi	Nation Party
EP	16	Emek Partisi	Labour Party
HÜDAPAR	15	Hür Dava Partisi	Free Cause Party
PS	14	Siyasi Semboller	Political Symbols
TDP (1)	12	Türkiye Değişim Partisi	Turkey Change Party
AO	11	Alperen Ocakları	Alperen Hearths
HEPAR	10	Hak ve Eşitlik Partisi	Rights and Equality Party
ANAP	9	Anavatan Partisi	Motherland Party
TDP (2)	8	Toplumcu Demokrasi Partisi	Socialist Democracy Party
TIP	5	Türkiye İşçi Partisi	Workers' Party of Turkey
ÖDP	5	Özgürlük ve Dayanışma Partisi	Freedom and Solidarity Party
TKH	5	Türkiye Komünist Hareketi	Communist Movement of Turkey
LDP	4	Liberal Demokrat Parti	Liberal Democratic Party
SDP	4	Sosyal Demokrasi Partisi	Social Democracy Party
DBP	4	Demokratik Bölgeler Partisi	Democratic Regions Party
GP	4	Gelecek Partisi	Future Party
BDP	3	Barış ve Demokrasi Partisi	Peace and Democracy Party
SCP	2	Sosyalist Cumhuriyet Partisi	Socialist Republican Party
HKP	2	Halkın Kurtuluş Partisi	People's Liberation Party
SYKP	1	Sosyalist Yeniden Kuruluş Partisi	Socialist Refoundation Party
SOL PARTI	1	Sol Parti	Left Party
KSP	1	Kürdistan Sosyalist Partisi	Kurdistan Socialist Party
HDP (2)	1	Hür Dava Partisi	Free Cause Party
TOTAL	18035		

Directorate of Mapping. Overlaying this polygon with a uniform 20 m×20 m grid produced a comprehensive set of query points compatible with Google's Street View sampling density. For each centroid we requested Street View metadata, yielding approximately 4.682 million candidate images. Restricting the selection to the months immediately surrounding local and general elections—when visual campaigning is most intense—reduced the corpus to 748,783 panoramas.[1] Each panorama carries a persistent reference ID, 360° horizontal and 180° vertical coverage in equirectangular projection, and precise latitude–longitude coordinates furnished by the API's geocoding services. Collecting the associated metadata and downloading the images required roughly six months of continuous processing on a 2019 MacBook Pro and resulted in a dataset of about 4 TB, underscoring the logistical demands of large-scale visual political analysis. A detailed breakdown of the available data is provided in Table 3.

Google Street View (GSV) offers unrivaled street-level coverage across Istanbul, yet that coverage is uneven, with gaps in narrow alleys, pedestrian-only zones, and restricted areas such as gated communities or military sites. These omissions arise from both physical access constraints and intentional exclusions for security or privacy, while infrequent image updates—including the pandemic-related pause—create temporal biases that can misrepresent rapidly changing features such as new infrastructure or informal street vendors. Nonetheless, GSV's breadth and public accessibility make it superior to any other visual data-collection system; researchers should simply interpret its outputs with caution when analyzing fine-grained or time-sensitive urban phenomena.

The polygonal boundaries of all Istanbul precincts (*mahalle*) were obtained as a single GeoJSON archive from the Istanbul Metropolitan Municipality's open-data portal and saved locally as `mahalle.geojson`. A companion CSV published by the Turkish Statistical Institute (TÜİK) supplies 2022 population totals, average household income, and detailed education shares for each precinct. Both resources are ingested by a custom Python script. The routine first invokes `GeoPandas` to read the GeoJSON; if the municipal export contains spurious trailing commas, a regex cleaner rewrites the file and retries the import. Geometries supplied in geographic coordinates (EPSG 4326) are re-projected to Web Mercator (EPSG 3857), after which the script computes each polygon's bounding box, centroid, and surface area and prints null-value diagnostics to validate the schema. Finally—performed in the notebook but omitted from the listing—the socio-economic CSV is left-joined on the precinct key, yielding a single `GeoDataFrame` that marries high-resolution boundaries to a rich attribute table suitable for regression analysis. Figure 3 displays the boundaries of all 1,009 *mahalles* of the Istanbul province, with each district highlighted in a corresponding color.

[1] The analysis covers all nationwide elections in Turkey since 2000. Municipal contests took place in March 2004, 2009, 2014, 2019, and 2024; parliamentary elections were held in November 2002, July 2007, June 2011, June 2015, the snap poll of November 2015, June 2018, and May 2023. For every contest, street-view imagery from *both* the election month and the immediately preceding month is examined.

Table 3: Monthly picture counts for election and pre-election months
Note. Any month not listed has zero available pictures for every district.

District	2011-05	2014-02	2014-03	2015-10	2015-11	2018-05	2018-06	2019-02	2019-03	2023-04	2023-05	2024-02	2024-03	Total
Adalar	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Arnavutköy	0	0	0	0	0	5	0	0	0	0	0	0	0	5
Ataşehir	0	8 926	27 271	4 667	0	2	55	0	0	54	0	0	0	40 975
Avcılar	0	0	0	0	0	0	0	0	0	0	0	0	78	78
Bağcılar	0	0	48	24 531	0	0	0	0	0	0	0	0	0	24 579
Bahçelievler	0	0	294	13 301	0	1	0	0	0	0	0	0	0	13 596
Bakırköy	0	0	151	7 380	0	5	1	0	0	0	0	0	0	7 537
Başakşehir	0	0	43	9 508	0	7	0	0	0	0	0	0	47	9 605
Bayrampaşa	0	0	177	1 170	0	0	0	0	0	0	0	208	0	1 555
Beşiktaş	0	0	10 617	0	0	0	0	0	0	0	0	0	0	10 617
Beykoz	0	41 348	11 126	1 179	0	5	0	1	0	40	0	0	0	53 699
Beylikdüzü	0	0	0	0	0	0	25	0	2	0	0	0	0	27
Beyoğlu	0	0	2 617	40	1	0	2	0	0	0	7	97	0	2 764
Büyükçekmece	0	0	0	0	0	1	0	0	0	0	0	0	0	1
Çatalca	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Çekmeköy	0	14 621	6 442	5 378	0	4	0	0	0	358	0	0	0	26 803
Esenler	0	0	0	4 487	0	0	0	0	0	0	45	0	0	4 532
Esenyurt	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Eyüpsultan	0	0	2 117	123	0	1	0	0	0	0	248	147	0	2 636
Fatih	0	0	95	9 111	0	9	10	0	1	0	0	358	0	9 584
Gaziosmanpaşa	0	0	258	0	0	2	0	0	0	0	836	165	0	1 261
Güngören	0	0	67	2 711	0	1	0	0	0	0	0	0	0	2 779
Kadıköy	1	14 349	3 410	12 901	0	6	1	0	0	0	0	0	0	30 667
Kağıthane	0	0	1 043	0	0	3	0	0	0	0	14	7	0	1 067
Kartal	0	17 226	20 145	0	0	1	0	0	0	195	0	13	0	37 580
Küçükçekmece	0	0	23 150	7 300	0	4	0	0	0	0	0	0	0	30 454
Maltepe	0	3 078	32 494	0	0	1	0	0	0	2	0	0	1	35 576
Pendik	0	44 350	21 243	0	0	1	0	0	0	3 521	0	0	0	69 115
Sancaktepe	0	5 804	32 849	1 381	0	0	21	0	0	1 023	0	0	0	41 078
Sarıyer	0	0	592	0	0	1	0	3	1	469	0	0	0	1 063
Silivri	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Şile	0	457	899	0	0	1	0	0	0	23	0	0	0	1 380
Sultanbeyli	0	8 678	32 926	0	0	2	5	0	0	3 562	0	0	0	45 173
Sultangazi	0	0	0	0	0	0	0	0	0	0	773	0	0	773
Şişli	0	0	5 442	0	0	1	0	0	4	0	0	29	0	5 476
Tuzla	0	631	24 925	0	0	1 777	3	0	4	6 765	0	0	0	26 303
Ümraniye	0	12 426	53 513	54 406	0	1	0	0	0	1 461	0	0	0	121 807
Üsküdar	0	22 185	23 508	35 571	0	1	0	3	0	0	0	0	0	81 268
Zeytinburnu	0	0	369	12 784	0	1	122	0	0	0	0	42	0	13 318
Total	1	196 438	328 492	208 986	1	75	242	1	12	11 420	1 923	1 066	126	748 783

The merged dataset includes (i) administrative identifiers such as unique feature IDs, province and district codes, and the precinct code itself; (ii) spatial metadata that includes official name, settlement type, postal code, distance to the district center, planimetric surface area, and multiple shape descriptors automatically generated by the GIS software; (iii) infrastructure proxies, including counts of registered households, total dwelling units, private workplaces, banks, and ATMs; and (iv) demographic detail comprising resident population with male/female splits, fifteen five-year age brackets, ten educational-attainment bins (from illiteracy through doctoral degrees), marital-status counts by gender, and

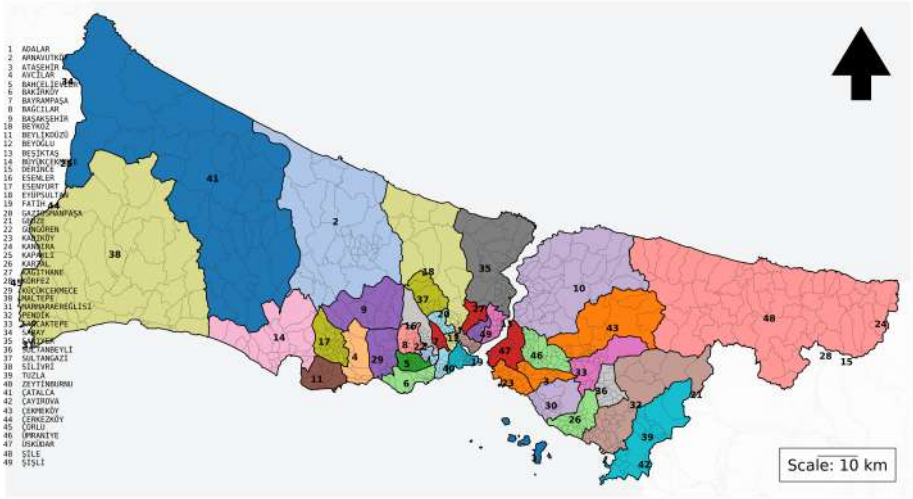

Fig. 3: District Boundries, Istanbul

precinct-level density and parent-district totals. This multi-dimensional portrait of Istanbul's *mahalles* provides the socio-economic and built-environment controls required for the study's empirical specifications.

6 Computer Vision Framework for Political-Symbol Detection

6.1 Two–Stage Object-Detection Framework

Detecting political insignia in 360° panoramas posed two main challenges: (i) symbol classes varied across election cycles, and (ii) street-level flags typically occupied only a few dozen pixels in the raw 16 000 × 8 000 images. To address these issues we designed a two–stage pipeline that combined open-vocabulary proposal generation with a vision transformer (OwlViT) and class-specific confirmation with a YOLOv12 model trained on manually annotated symbols. The computational footprint is summarized in Table 4.

Stage 1: Zero-Shot Proposal Generation with OWL-ViT We adopted the DINO-OWL variant of **OwlViT** (Minderer et al., 2022; Yin et al., 2023), a masked Vision Transformer fine-tuned for open-vocabulary object detection. Given an image and a text prompt, the model returned bounding boxes $B_k = (x_1, y_1, x_2, y_2)$ with confidence scores $\hat{p}_k \in [0, 1]$. Because OWL-ViT was trained on millions of image–text pairs, it generalizes well to unseen iconography and requires no flag-specific tuning.

Pre-processing. Each panorama was (i) converted to RGB and cropped to remove sky and ground borders, (ii) sliced into a 4 × 2 grid of eight overlapping tiles (about 2 000 × 2 000 px each), and (iii) passed through OWL-ViT with the query

string "a flag or a political symbol" at a confidence threshold of 0.10. Candidate boxes were written to an `.xlsx` file together with run-time diagnostics. A three-way batching scheme made it easy to parallelize processing across GPUs.

Stage 2: Class-Specific Confirmation with YOLO12s The second stage re-classified every OWL-ViT proposal with a lightweight YOLO12s detector (C.-Y. Wang et al., 2024). YOLO's decoupled head and dilated-encoder backbone delivered state-of-the-art accuracy on objects smaller than 128 px while maintaining real-time throughput on commodity GPUs.

We manually annotated 18,035 party symbols across 6 496 unique panoramas in Roboflow, achieving strong inter-annotator agreement (Fleiss Kappa = 0.82). The annotated set was exported to the Ultralytics format and divided into training, validation, and test splits in an 80/10/10 ratio. Model training ran for 100 epochs on 640-pixel images with a batch size of 32, used early stopping with a patience of 20 epochs, relied on mixed-precision computation (`half=True`), and employed four data-loader workers.

Model selection. We benchmarked variants v8–v12; YOLO12s obtained the best scores for our large annotation set (all classes (v40)). To maximize recall on extremely small targets, we applied two preprocessing strategies. First, each OWL-ViT tile underwent secondary slicing with SAHI's 0.3-overlap rule, which ensured that the smallest flags ($\approx$ 40 px in the full panorama) occupied at least a 64×64 px footprint in one of the resulting crops and increased recall by roughly six percentage points. Second, we removed non-informative image borders—about 14 % of all pixels—which lowered Stage 1 computational cost by the same margin. A summary of the workload is provided in Table 4.

Table 4: Workload summary for the Istanbul case study.

Quantity	**Value**
Candidate panoramas queried (GSV metadata)	4 682 647
Panoramas in election / pre-election months	748 783
Panoramas within 20 m of party offices	379 158
Unique annotated panoramas (deduplicated)	6 496
Manual symbol annotations	18 035
Street-level party-office addresses	14 583
OWL-ViT forward passes	6 000 000
YOLO inference crops	25 000 000

6.2 Benchmarking YOLO v8–v12 on Turkish Political-Symbol Detection

All five YOLO variants (v8–v12) were fine-tuned under the same training regimen: 100 epochs, image resolution 640 px, batch size 32, mixed-precision computations (`half=True`), early stopping with a patience of 20 epochs, and four data-loader workers. Backbone weights were initialized from their official Ultralytics checkpoints; only the detection heads were re-initialized.

Two curated datasets were used, and finally only one of them (v40) was used in the analysis due to its more generalizable performance. *Version 40 (v40)* contained 6 496 panoramas and 18,035 base instances spanning the Turkish flag, Atatürk iconography, the five parliamentary parties (AKP, CHP, MHP, HDP, and İYİ), and several minor parties preserved from earlier work. Its training split was expanded offline with a five-fold augmentation pipeline—horizontal flips, up-to-30% random zoom crops, 15 grayscale, and ±15% brightness jitter—yielding 68,680 augmented training images and 72,979 images in total. *Version 41 (v41)* also comprised 6,496 panoramas but 22,917 base instances; it focused on the five major parties plus the Turkish flag and Atatürk symbols, dropping all minor parties. The training split received a three-fold Roboflow augmentation with the same transform suite, producing 41,208 augmented training images and 45,507 images overall.

Models were evaluated on a held-out 20% split of each dataset with COCO metrics (mAP_{50}, $mAP_{50:95}$, precision, recall). For every run the checkpoint with the highest $mAP_{50:95}$ is reported in Table 5.

Table 5: YOLO v8–v12 results on v40 and v41; bold indicates the top performer per dataset.

Model	Dataset	Epoch	Time†	mAP_{50}	$mAP_{50:95}$	Prec.	Rec.
YOLO8	v41	41	4.35	0.836	0.836	0.847	0.775
YOLO9	v41	19	4.99	0.831	0.831	0.822	0.792
YOLO10	v41	20	2.84	0.837	0.824	0.817	0.800
YOLO11	v41	19	2.76	0.837	0.837	0.830	0.776
YOLO12	**v41**	**53**	**14.16**	**0.838**	**0.838**	**0.861**	**0.786**
YOLO8	v40	65	15.12	0.715	0.715	0.842	0.654
YOLO9	v40	74	36.48	0.740	0.740	0.867	0.677
YOLO10	v40	31	8.62	0.703	0.703	0.772	0.677
YOLO11	v40	37	9.26	0.722	0.719	0.832	0.645
YOLO12	**v40**	**66**	**20.64**	**0.750**	**0.747**	**0.750**	**0.695**

†Hours to best epoch, rounded to two decimals.

YOLO12 trained on v41 produced the highest overall score ($mAP_{50:95}$ = 0.838), exceeding the next-best contender (YOLO11 on v41) by roughly two percentage points and its own v40 analog by more than twelve. A key driver of this jump is the narrower label space in v41: after minor-party symbols were removed, the model needed to activate fewer output channels. The resulting reduction in inter-class confusion raised precision across the board, yielding an average $mAP_{50:95}$ gain of 11.8 percentage points over v40. Although YOLO11 converged five times faster (about 2.8 hours) and lagged YOLO12 by only 0.1 percentage point in precision, it recovered fewer true positives and therefore trailed in recall. On v40, YOLO9 delivered the highest precision (0.867) but at the cost of markedly lower recall, a trade-off acceptable only when false positives are especially costly. For all subsequent empirical analyses two (best) checkpoints were frozen:

– **v40** – YOLO12, epoch 66 ($mAP_{50:95}$ = 0.747, precision 0.750, recall 0.695);

– **v41** – YOLO12, epoch 53 ($\text{mAP}_{50:95} = 0.838$, precision 0.861, recall 0.786).

In Figures 4 and 5, we provide a short overview of the results produced by our data-acquisition and object-detection pipeline, including Google Street View coverage, detected political party symbols, and recorded depictions of Atatürk and the Turkish flag for metropolitan Istanbul and four representative districts—Sultanbeyli and Ümraniye (traditionally conservative) and Beşiktaş and Kadıköy (traditionally liberal). Detailed detection maps for every district in the province are provided in the Appendix.

7 Empirical Approach

Our empirical strategy unfolds in four steps, each matching one of the paper's central hypotheses. We construct a monthly, district–party panel that merges geolocated flag detections from Google Street View with official election returns and a rich set of socioeconomic indicators (Table 6). Every analysis is replicated under three image-classifier confidence thresholds—RAW, $c > 0.90$, and $c > 0.95$—and all models absorb variation in Street-View sampling intensity via the panorama count O_{it}. Heteroskedasticity-robust standard errors are clustered by district.

Variable(s)	Definition
`YAS_T_15_1`, `YAS_T_20_2`, `YAS_T_65_A`	Shares of residents aged 15–20, 20–25, and 65+
`EG_LISANS`	Proportion of adults with at least a bachelor's degree
`KADIN_NUFU`	Female share of the population
`KONUTTOPLA`	Number of residential units (housing stock)
`OZELISYERI`	Number of private workplaces (commercial density)

Table 6: Precinct-level socioeconomic covariates used in all H1–H4 specifications.

7.1 H1 Symbol Density and Electoral Outcomes

H1a Symbol Density → Vote Share For party p in district i and month t we regress the official vote share

$$\text{VoteShare}_{pit} \;=\; \frac{\text{ballots for } p}{\text{valid ballots}_{it}} \times 100$$

on the flag count S^c_{pit} (confidence cut c) and coverage O_{it}. Three nested OLS models are estimated:

$$\begin{aligned}
\text{VoteShare}_{pit} &= \alpha + \beta S^c_{pit} + \gamma O_{it} + u_{pit} && \text{(baseline)}\\
&= \alpha + \beta S^c_{pit} + \gamma O_{it} + \delta_p + \lambda_t + \mu_i + u_{pit} && (+ \text{ party, month, district FE})\\
&= \alpha + \beta S^c_{pit} + \gamma O_{it} + \mathbf{X}_i'\boldsymbol{\theta} + \delta_p + \lambda_t + \mu_i + u_{pit} && (+ \text{ socioeconomic controls}).
\end{aligned}$$

A positive and significant $\hat{\beta}$ across tiers supports H1a.

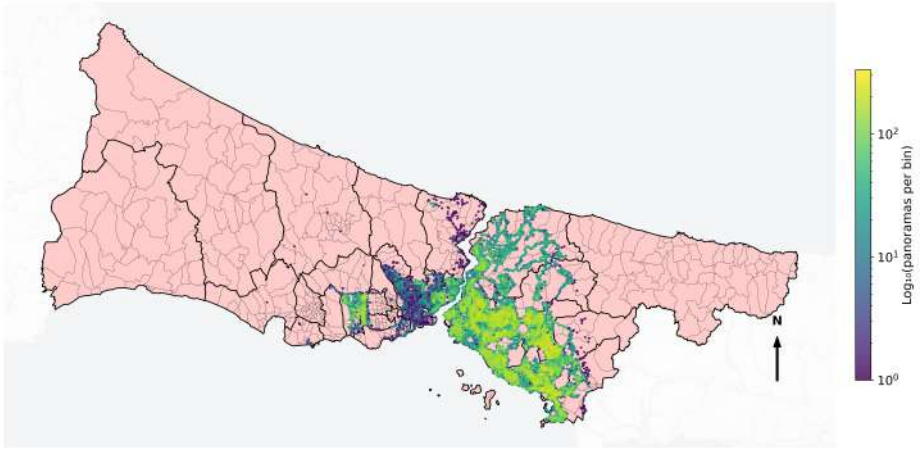

(a) Available Google Street View panoramas

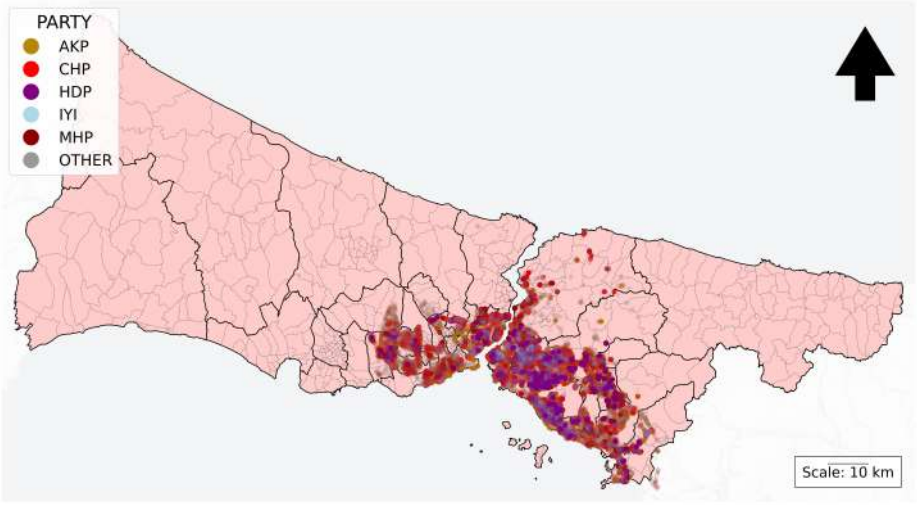

(b) Party symbols

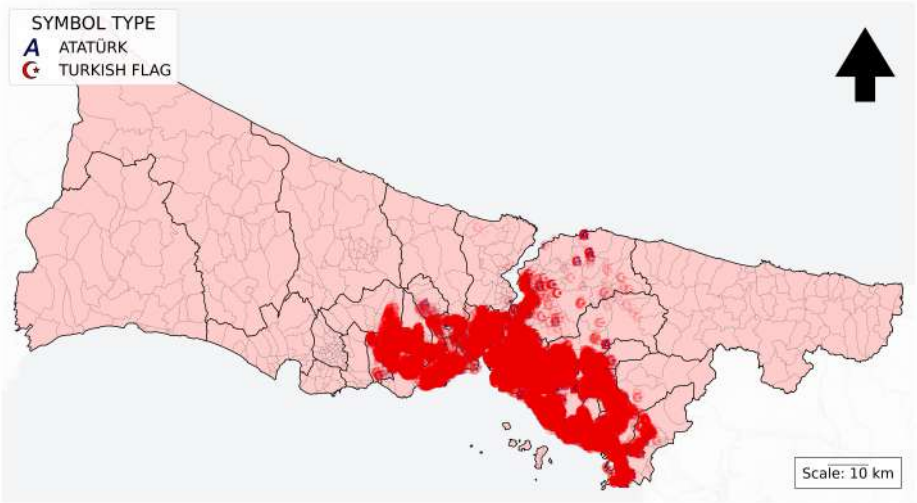

(c) Atatürk and Turkish flag

Fig. 4: Symbol detection results from Istanbul

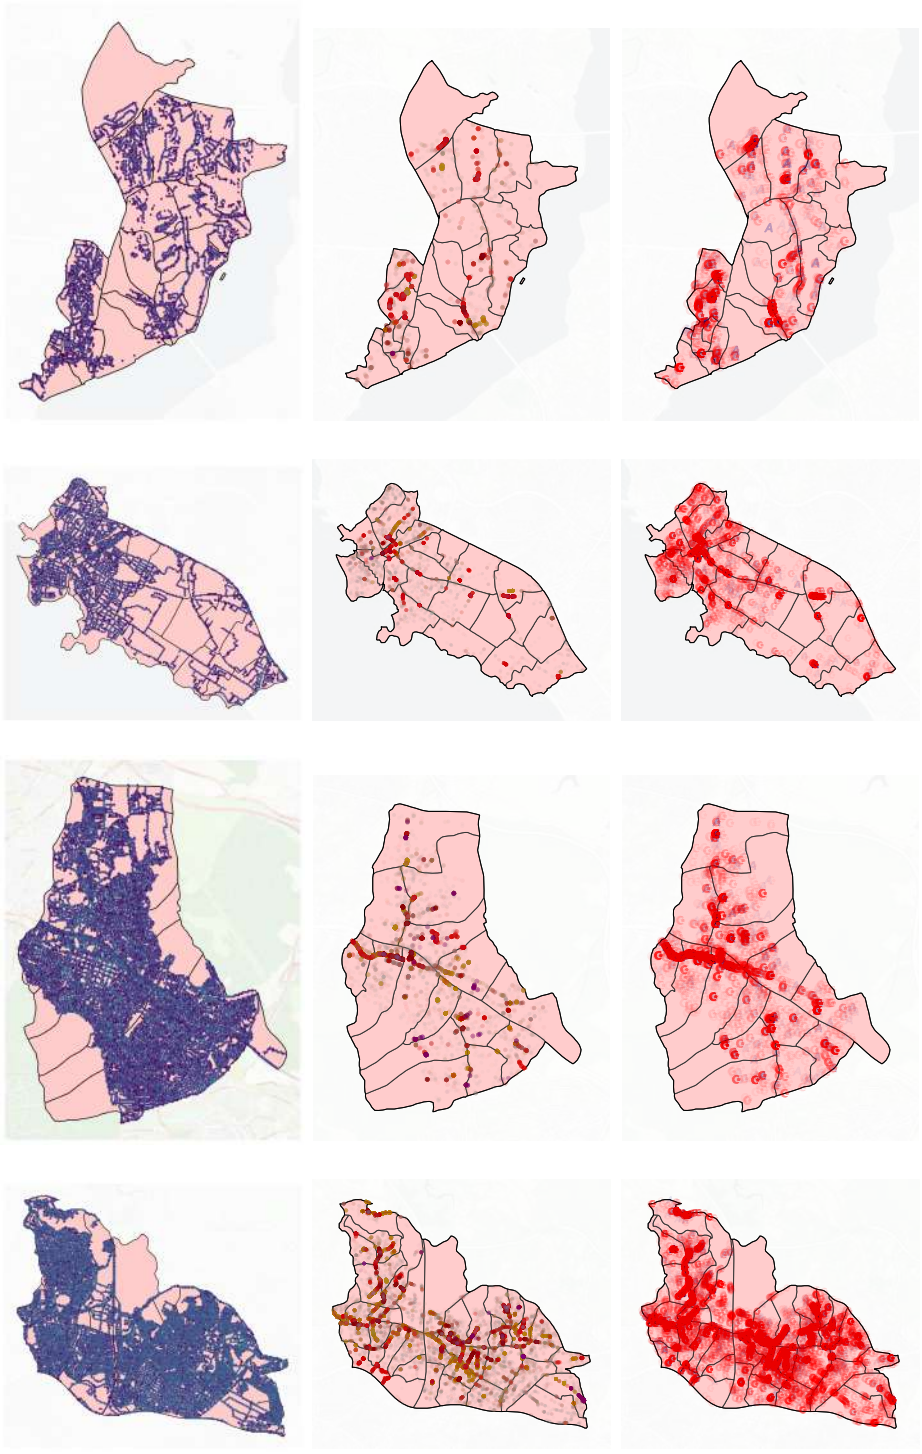

Fig. 5: Illustration of Google Street View panorama coverage (Column 1), political party symbol detections (Column 2), and occurrences of Atatürk and the Turkish flag (Column 3) across four Istanbul districts: Beşiktaş (Row 1), Kadıköy (Row 2), Sultanbeyli (Row 3), and Ümraniye (Row 4).

H1b Symbol Density → Turnout Replacing the party-level outcome with district turnout,

$$\text{Turnout}_{it} \;=\; \frac{\text{ballots cast}_{it}}{\text{eligible voters}_{it}} \times 100,$$

we aggregate flags to $S^c_{it} = \sum_p S^c_{pit}$ and re-estimate the same model sequence. A significant, positive $\hat{\beta}$ indicates that visual saturation mobilizes voters.

7.2 H2 Symbolic Competition and Coordination

For every district–month (i, t) we summarize visual competition with the Herfindahl-based flag-diversity index

$$\mathcal{D}_{it} = 1 - \sum_p \left(\frac{F_{pit}}{\sum_q F_{qit}} \right)^2,$$

which equals 0 when a single party supplies all detected symbols and approaches 1 as flags are distributed more evenly across parties. Electoral dominance is captured by the winner–runner-up vote-share gap

$$\text{Margin}_{it} = v^{(1)}_{it} - v^{(2)}_{it}.$$

We regress Margin_{it} on $\mathcal{D}_{it}$, controlling for turnout, population density, month fixed effects, and—when indicated—the full socioeconomic block listed in Table 6. A negative coefficient on $\mathcal{D}_{it}$ ($\hat{\beta} < 0$) supports H2, implying that greater visual pluralism coincides with narrower victory margins.

7.3 H3 Spatial Strategies

H3a Precinct-Level Spatial Avoidance. Within each precinct we form all unordered party pairs (A, B) and model the mean inter-flag distance

$$d_{ABj} = \alpha + \delta\, 1_{\text{Ally}_{AB}} + \eta_1 D_{Aj} + \eta_2 D_{Bj} + \mathbf{X}'_j \boldsymbol{\theta} + \mu_i + \lambda_t + \varepsilon_{ABj}.$$

A negative $\hat{\delta}$ implies that allied parties place their flags closer together than rivals, thus supporting H3a.

H3b Incumbent vs. Challenger Spatial Strategy. For every party we calculate two normalized dispersion metrics—convex-hull area per flag and mean nearest-neighbor distance per flag—and compare incumbents (AKP–MHP) to challengers via Welch t-tests. Evidence that incumbents display tighter clustering validates the defensive-clustering/offensive-dispersion hypothesis.

7.4 H4 Demographic and Built-Environment Correlates

At the precinct level we regress party-specific flag counts Sym_{ip} on each socioeconomic attribute in Table 6:

$$\mathrm{Sym}_{ip}^{(c)} = \alpha_{pj}^{(c)} + \beta_j^{(c)} X_{ij} + \varepsilon_{ipj}^{(c)}, \qquad j \in \{1, \ldots, 7\},\ c \in \{\mathrm{raw}, 0.90, 0.95\}.$$

Significant coefficients reveal how parties target flags toward specific demographic or built-environment profiles, completing the behavioral narrative established by H1–H3.

Across all hypotheses we (i) use identical confidence cuts, (ii) control for panorama coverage, and (iii) cluster heteroskedasticity-robust errors by district. These conventions guarantee that findings from H1 through H4 are directly comparable and collectively illuminate the spatial logic of party symbolism in Istanbul.

8 Results

H1a. Across all nine specifications, symbol density remains a statistically robust and substantively meaningful predictor of district-level vote share, although the magnitude of the association diminishes once unobserved heterogeneity is absorbed by fixed effects. In the baseline models (columns 1–3), which adjust only for the number of panoramas sampled, an additional 10,000 detected flags corresponds to an 8.1-percentage-point increase in a party's vote share when all detections are included ($\beta = 0.081$, SE $= 0.024$, $p < .01$). The estimated effect rises to 15.7 and 18.4 points, respectively, when the sample is restricted to detections with confidence scores above 0.90 and 0.95. Introducing party, date, and district fixed effects (columns 4–6) reduces these coefficients to 1.5, 3.4, and 4.1 percentage points, implying that roughly four-fifths of the raw association is explained by time-invariant party characteristics and stable spatial or temporal factors. Adding the full block of fourteen socioeconomic controls (columns 7–9) leaves the point estimates and their standard errors virtually unchanged, underscoring the robustness of the fixed-effects results to demographic and built-environment adjustments. Model fit improves dramatically with the inclusion of fixed effects, as R^2 increases from approximately 0.076 to 0.836, confirming that most of the explained variance is driven by interparty, interdistrict, and intertemporal differences rather than by observable precinct characteristics. Standard errors, clustered at the district level, indicate that all reported coefficients remain significant at the 5 percent level or better, lending credence to the inference that a denser symbolic presence in the street translates into measurably higher electoral returns for the sponsoring party. The results are shown in Table 7.

Table 7: H1a – OLS estimates of symbol density on party vote share

	Baseline			**Fixed-effects only**		
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	0.195***	0.195***	0.195***	0.069***	0.069***	0.069***
	(0.008)	(0.008)	(0.008)	(0.006)	(0.006)	(0.006)
Symbol Count (All) ($\times 10^4$)	0.081***			0.015***		
	(0.024)			(0.006)		
Symbol Count (Conf. >0.90) ($\times 10^4$)		0.157***			0.034**	
		(0.038)			(0.013)	
Symbol Count (Conf. >0.95) ($\times 10^4$)			0.184***			0.041**
			(0.045)			(0.016)
Party fixed effects				X	X	X
Date fixed effects				X	X	X
District fixed effects				X	X	X
Panorama-count control	X	X	X	X	X	X
R^2	0.077	0.076	0.075	0.835	0.836	0.836
Adj. R^2	0.072	0.071	0.070	0.810	0.810	0.810
N	357	357	357	357	357	357

	Social Variables + Fixed Effects		
	(7)	(8)	(9)
Intercept	0.054***	0.054***	0.054***
	(0.004)	(0.004)	(0.004)
Symbol Count (All) ($\times 10^4$)	0.015***		
	(0.006)		
Symbol Count (Conf. >0.90) ($\times 10^4$)		0.034**	
		(0.014)	
Symbol Count (Conf. >0.95) ($\times 10^4$)			0.041**
			(0.016)
Party fixed effects	X	X	X
Date fixed effects	X	X	X
District fixed effects	X	X	X
Social-control block[a]	X	X	X
Panorama-count control	X	X	X
R^2	0.835	0.836	0.836
Adj. R^2	0.810	0.810	0.810
N	357	357	357

Notes: District-clustered robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

[a] Social-control variables: `YUZOLCUMU` (area in km^2), `SOKAKSAYIS` (number of streets), `HANE_SAYIS` (number of households), `OZELISYERI` (private workplaces), `BANKA` (bank branches), `ATM` (ATMs), `KADIN_NUFU` (female population), `YOGUNLUK` (population density), `ILCE_NUFUS` (district population), `EG_LISANS` (bachelor's degree attainment), `HIcEVLENME` (never-married share), `YAS_T_15_1` (age 15–20), `YAS_T_20_2` (age 20–25), `YAS_T_65_A` (age 65+).

H1b. At the district–month level, symbol density exhibits no substantively meaningful or statistically reliable association with voter turnout. In the parsimonious baseline specifications (columns 1–3), the estimated effect of an additional 10,000 detected symbols on turnout is extremely small—0.0006 in column 1, 0.0017 in column 2, and 0.0020 in column 3—equivalent to increases of just 0.06 to 0.20 percentage points; none of these coefficients approach conventional significance thresholds. Once month and district fixed effects are introduced (columns 4–6), the point estimates not only lose what little magnitude they had but also turn slightly negative (–0.0000 to –0.0002), again with p-values far from significance. Adding the full block of fourteen socioeconomic controls (columns 7–9) leaves the coefficients effectively unchanged and still statistically indistinguishable from zero. Model fit rises dramatically with the inclusion of fixed effects (R^2 increases from roughly 0.045 to 0.993), indicating that nearly all systematic variation in turnout is captured by stable temporal and spatial factors rather than by symbolic saturation in the streets. Clustered standard errors confirm that none of the symbol-density coefficients are significant even at the 10 percent level. In short, while denser flag displays are linked to higher vote shares for the sponsoring party, they do not translate into greater aggregate voter participation. The results are provided in Table 8.

H2a. Across all nine specifications we find no convincing evidence that district-level *symbol diversity*—captured with the Herfindahl-based index $\mathcal{D}_{it} = 1 - \sum_p \left(F_{pit} / \sum_q F_{qit}\right)^2$, which equals 0 when one party monopolizes all flags and approaches 1 when symbols are evenly shared—systematically influences the victor's margin of victory. In the baseline models that control only for population density and turnout (columns 1–3), an additional increase in symbol diversity is associated with 14.5–33.5 fewer votes separating the first- and second-place parties, but the attendant standard errors (14.4–52.3) render these estimates statistically indistinguishable from zero. Incorporating month fixed effects (columns 4–6) roughly doubles the absolute size of the coefficients (–24.7 to –55.9) and raises R^2 from about 0.11 to 0.26, yet the estimates remain imprecise and the adjusted R^2 stays near zero, underscoring that fixed temporal shocks, not flag diversity, explain most of the incremental fit. Once both month effects and the full battery of socioeconomic controls are introduced (columns 7–9), the point estimates on symbol diversity shrink toward the origin, even switching sign in the highest-confidence specification (+2.0), while their standard errors inflate four- to tenfold. Although these saturated models capture roughly 85% of the total variance, the corresponding adjusted R^2 values (0.34–0.42) indicate considerable overfitting given the very small sample ($N = 28$). Coefficients on turnout are consistently negative (–74.9 to –501.8) but likewise never attain statistical significance, and population-density effects hover around zero throughout. Overall, the data provide little support for Hypothesis 2a: variation in the diversity of visible partisan symbols does not translate into meaningful or reliable shifts in electoral victory margins once temporal, geographic, and demographic factors are taken into account. The estimates are given in Table 9.

Table 8: H1b – OLS estimates of symbol density on voter turnout

	Baseline			**Fixed effects only**		
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	0.879***	0.879***	0.879***	0.905***	0.905***	0.906***
	(0.0051)	(0.0051)	(0.0051)	(0.0026)	(0.0026)	(0.0026)
Symbol count ($\times 10^4$)	0.0006			−0.0000		
	(0.0016)			(0.0002)		
Symbol count > 0.90 ($\times 10^4$)		0.0017			−0.0001	
		(0.0043)			(0.0007)	
Symbol count > 0.95 ($\times 10^4$)			0.0020			−0.0002
			(0.0050)			(0.0008)
Month fixed effects				X	X	X
District fixed effects				X	X	X
Social-control block[a]						
Observation-count control	X	X	X	X	X	X
R^2	0.045	0.045	0.045	0.993	0.993	0.993
Adj. R^2	0.020	0.020	0.020	0.988	0.988	0.988
N	81	81	81	81	81	81

	Social controls + FE		
	(7)	(8)	(9)
Intercept	0.580***	0.580***	0.580***
	(0.0018)	(0.0018)	(0.0018)
Symbol count ($\times 10^4$)	−0.0000		
	(0.0002)		
Symbol count > 0.90 ($\times 10^4$)		−0.0001	
		(0.0007)	
Symbol count > 0.95 ($\times 10^4$)			−0.0002
			(0.0008)
Month fixed effects	X	X	X
District fixed effects	X	X	X
Social-control block[a]	X	X	X
Observation-count control	X	X	X
R^2	0.993	0.993	0.993
Adj. R^2	0.988	0.988	0.988
N	81	81	81

Notes: District-clustered robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

[a] Social-control variables: `YUZOLCUMU` (area in km^2), `SOKAKSAYIS` (number of streets), `HANE_SAYIS` (number of households), `OZELISYERI` (private workplaces), `BANKA` (bank branches), `ATM` (ATMs), `KADIN_NUFU` (female population), `YOGUNLUK` (population density), `ILCE_NUFUS` (district population), `EG_LISANS` (bachelor's degree attainment), `HIcEVLENME` (never-married share), `YAS_T_15_1` (age 15–20), `YAS_T_20_2` (age 20–25), `YAS_T_65_A` (age 65+).

Table 9: H2a — OLS estimates of symbol diversity on victory margin.

	Baseline controls only			**+ Month fixed effects**			**+ Month FE & Socio controls**		
	Raw	≥ 0.90	≥ 0.95	Raw	≥ 0.90	≥ 0.95	Raw	≥ 0.90	≥ 0.95
Intercept	321.990	398.006	384.462	492.193	475.151	458.432	221.742	96.412	142.570
	(255.779)	(277.765)	(289.087)	(337.591)	(314.578)	(307.092)	(826.497)	(1007.840)	(757.598)
Symbol diversity	–33.531	–22.182	–14.534	–55.919	–30.661	–24.730	–13.396	–1.098	1.958
	(52.336)	(16.222)	(14.445)	(39.234)	(26.735)	(15.646)	(146.764)	(151.245)	(53.537)
Turnout (%)	–328.956	–410.640	–398.387	–501.757	–478.759	–463.024	–222.102	–74.926	–128.233
	(279.258)	(303.165)	(315.193)	(370.073)	(343.762)	(337.397)	(897.486)	(1079.752)	(826.012)
Pop. density	0.000008	–0.000105	–0.000097	–0.000022	–0.000083	–0.000041	–0.000041	0.000058	–0.000140
	(0.000181)	(0.000196)	(0.000213)	(0.000197)	(0.000234)	(0.000211)	(0.000619)	(0.001131)	(0.000838)
Month FEs	–	–	–	X	X	X	X	X	X
Socio controls	–	–	–	–	–	–	X	X	X
R^2	0.095	0.126	0.110	0.200	0.262	0.241	0.843	0.854	0.850
Adj. R^2	–0.019	0.017	–0.001	–0.029	0.004	0.024	0.395	0.344	0.422
Observations	28	28	28	28	28	28	28	28	28

H3a. The regression evaluates whether formal electoral alliances or the sheer abundance of each party's flags predict the physical proximity of their symbols at the precinct level. The intercept implies that, on average, two *non-allied* parties place their flags roughly 622 meters apart (see column 1; SE $\approx$ 1.4 m). By contrast, the coefficient on *Allied_AB* is statistically different from zero but substantively negligible: the point estimate is smaller than one millimeter in every specification (reported as –0.000 after rounding), indicating that even a shift from a non-allied to an allied pairing leaves the expected separation virtually unchanged.

Flag abundance likewise exerts a trivial influence. An extra 10,000 detections for party A (D_{A}) reduces the mean inter-party distance by only 1–4 centimeters, while the same increment for party B (D_{B}) shortens the gap by roughly 3–13 centimeters; all estimates are highly significant owing to the enormous sample size (N up to 102,672 pairs) yet remain far too small to matter in practice.

Including high-dimensional fixed effects for each party and for every district–month combination lifts the R^2 from essentially zero (< 0.01) to about 0.15, confirming that most explainable variation in flag spacing is tied to stable spatial-temporal factors rather than to alliance status or symbolic density. Adding fourteen socioeconomic covariates leaves fit metrics unchanged, reinforcing the conclusion that neither formal coalitions nor local demographics meaningfully govern how closely rival parties co-locate their visual propaganda. In short, the data offer no substantive support for the hypothesis that allied parties coordinate their flag placement to appear in closer physical proximity. The results can be seen in Table 10.

H3b. The dispersion pipeline operationalizes each party's spatial footprint as the convex-hull area spanned by its flags and then rescales that area by the number of symbols detected, yielding *hull_km*^{2} *per symbol*. A two-sample t-test compares this metric between the governing bloc (AKP and MHP) and all challenger parties.

Table 10: H3a – OLS estimates of symbol density on seat margins

	Baseline			**+ Fixed effects**			**+ FE & Social**		
	Raw	0.90	0.95	Raw	0.90	0.95	Raw	0.90	0.95
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Intercept	621.584***	683.169***	714.847***	564.806***	573.953*	860.716***	564.806***	573.953*	860.716***
	(1.417)	(4.054)	(5.495)	(140.30)	(332.36)	(36.30)	(140.30)	(332.36)	(36.30)
Allied_ AB	−0.000***	−0.000***	−0.000***	0.000***	0.000**	0.000	0.000***	0.000**	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
D_{A}	−0.014***	−0.032***	−0.043***	−0.016***	−0.029***	−0.040***	−0.016***	−0.029***	−0.040***
	(0.002)	(0.005)	(0.007)	(0.002)	(0.005)	(0.007)	(0.002)	(0.005)	(0.007)
D_{B}	−0.032***	−0.106***	−0.133***	−0.059***	−0.112***	−0.123***	−0.059***	−0.112***	−0.123***
	(0.003)	(0.013)	(0.016)	(0.003)	(0.013)	(0.015)	(0.003)	(0.013)	(0.015)
Party FE (A & B)				X	X	X	X	X	X
District × Month FE				X	X	X	X	X	X
Social-control block[a]							X	X	X
R^2	0.001	0.005	0.008	0.155	0.152	0.152	0.155	0.152	0.152
Adj. R^2	0.001	0.005	0.008	0.154	0.147	0.145	0.154	0.147	0.145
N	102 672	15 790	10 116	102 672	15 790	10 116	102 672	15 790	10 116

Notes: Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
[a] Social-control block contains demographic and built-environment covariates: `YUZOLCUMU` (area in km^2), `SOKAKSAYIS` (number of streets), `HANE_SAYIS` (number of households), `OZELISYERI` (private workplaces), `BANKA` (bank branches), `ATM` (ATMs), `KADIN_NUFU` (female population), `YOGUNLUK` (population density), `ILCE_NUFUS` (district population), `EG_LISANS` (bachelor's degree attainment), `HIcEVLENME` (never-married share), `YAS_T_15_1` (age 15–20), `YAS_T_20_2` (age 20–25), `YAS_T_65_A` (age 65+). Full sets of party and district×month dummies are included where indicated.

Across all three confidence cuts the results strongly corroborate the "defensive clustering versus offensive dispersion" hypothesis. Using every detection in the corpus (*raw* run, N = 1,217,695), challenger flags occupy significantly more territory per symbol than incumbent flags ($t = 3.652$, $p = 0.0008$). The pattern persists when the sample is restricted to high-confidence detections: with a 0.90 threshold (N = 454,392) the gap remains large and significant ($t = 2.952$, $p = 0.0079$), and the 0.95 threshold (N = 389,914) produces an almost identical test statistic ($t = 3.062$, $p = 0.0079$). Because the test statistic is positive, the mean convex-hull area per challenger symbol exceeds that of incumbents in every specification—exactly the pattern expected if incumbents "defensively" saturate their strongholds while challengers "offensively" disperse symbols to broadcast a wider territorial claim. The consistency of the finding across confidence cuts, coupled with p-values well below the conventional 1 percent level, provides robust empirical support for Hypothesis 3b.

H3b. To gauge how parties deploy visual resources, we compute for each organization the convex-hull area enclosing all of its detected flags and divide this surface by the number of symbols, yielding km^2 per symbol. A two-sample t-test then contrasts this dispersion metric for the governing bloc (AKP + MHP) with that for all challengers.

Across all three confidence thresholds the results are consistent with the "defensive-clustering versus offensive-dispersion" thesis. Using the full set of detections (*raw* run, N = 1,217,695), challenger symbols cover significantly more territory per flag than government symbols ($t = 3.652$, $p = 0.0008$). The dif-

ference remains significant when the analysis is limited to high-confidence flags: at the 0.90 cutoff ($N = 454{,}392$) we obtain $t = 2.952$ ($p = 0.0079$), and at the 0.95 cutoff ($N = 389{,}914$) $t = 3.062$ ($p = 0.0079$). In every specification the mean convex-hull area per challenger flag exceeds that of incumbents, implying that incumbents concentrate flags within strongholds while challengers spread a thinner stock across the metropolitan landscape.

Descriptive diagnostics reinforce this pattern. Table 11 reports that the AKP averages 0.007 km^2 per flag and a mean nearest-neighbor gap of only 2.1×10^{-3} m; the MHP, though smaller, still clusters tightly (0.047 km^2, 0.17 m). The principal opposition party, CHP, matches the government's density (0.010 km^2, 0.02 m), yet its ally İYİ disperses far more broadly (0.190 km^2, 2.67 m). Averaging by bloc, government parties occupy 0.027 km^2 per symbol, whereas opposition parties require 0.100 km^2—a nearly four-fold gap. Minor or newly founded parties (e.g., Alperen Ocakları, TKH) are even more diffuse, with per-symbol footprints in double-digit square kilometers.

Clustering appears driven less by incumbency than by organizational capacity. Parties able to finance and coordinate hundreds of thousands of banners—whether in power (AKP) or out (CHP)—blanket their core precincts, producing a "wall-to-wall" visual field. Resource-constrained actors, regardless of coalition status, adopt an extensive strategy: they scatter limited flags across a larger area to project breadth of presence. Thus, while the government–opposition contrast remains statistically clear, the underlying mechanism is better described as *capacity-based clustering*: high-capacity parties defend territory through saturation; low-capacity parties pursue reach through dispersion.

H4a. The regression evidence points to a single, robust demographic correlate of symbolic saturation. Precincts with comparatively larger proportions of young adults aged 20 to 25 exhibit systematically higher flag densities, and this relationship holds after controlling for party, district, and month fixed effects as well as the full suite of socioeconomic covariates. By contrast, other indicators—educational attainment, female population share, housing stock, commercial density, and overall population density—show no consistent or meaningful association with the distribution of party symbols. These patterns imply that parties strategically concentrate visual propaganda in areas where younger voters are most prevalent, while treating the remaining social characteristics as essentially neutral in their flag-deployment calculus. The results are displayed in Figure 6, where colored stars indicate statistically significant differences among the major parties. All pairwise regression tables are provided in the Appendix.

9 Discussion

The evidence assembled in this paper demonstrates that Istanbul's public streetscape constitutes a measurable—and electorally consequential—layer of political communication. Processing 748,783 election-month panoramas with an open-vocabulary proposal network followed by a domain-specific YOLOv12 classifier produced a

Table 11: Spatial dispersion by party and bloc

Party	**Bloc**	n	**km²/symbol**	**NNm/symbol**
Alperen Ocakları	Opposition	10	41.232110	355.861947
Emek Partisi	Opposition	932	2.504961	0.342106
HÜDA PAR	Government	2 965	0.836943	0.042211
AKP	Government	401 672	0.007236	0.000005
ANAP	Opposition	16	18.372837	223.277136
Atatürk symbols†	Neutral	31 474	0.081010	0.000743
BBP	Government	5 444	0.412748	0.010201
BDP	Opposition	219	9.459783	3.697884
BTP	Opposition	611	2.864477	0.667605
CHP	Opposition	254 265	0.010353	0.000015
DBP	Opposition	1 529	1.317078	0.139369
Demokrat Parti	Opposition	759	2.705787	0.383349
DEVA	Opposition	2 533	0.946846	0.055018
DSP	Government	10 299	0.211312	0.002641
Gelecek Partisi	Opposition	135	10.533819	7.994860
Hak ve Eşitlik Partisi	Opposition	9	27.067878	652.421879
HDP	Opposition	21 371	0.129911	0.000908
HKP	Opposition	4	95.300300	4 615.294236
Hür Dava Partisi	Government	7	38.741829	922.336114
HYP	Opposition	287	7.011158	2.559598
İşçi Partisi	Opposition	2 926	0.783046	0.044103
İYİ	Opposition	14 618	0.190092	0.002666
Memleket Partisi	Opposition	4 373	0.522723	0.013893
MHP	Government	59 303	0.047438	0.000168
Millet Partisi	Opposition	9 964	0.228625	0.003068
Özgürlük ve Dayanışma Partisi	Opposition	761	2.551762	0.384438
Saadet	Opposition	120 475	0.023915	0.000054
SDP	Opposition	14	52.671634	393.493543
TDP (Toplumcu Demokrasi)	Opposition	3 200	0.709508	0.031052
TİP	Opposition	4	31.453291	3 316.310395
TKH	Opposition	25	39.371697	144.460473
TKP	Opposition	157	8.511314	5.696202
Turkish flag†	Neutral	241 888	0.012323	0.000027
Türkiye Değişim Partisi	Opposition	830	2.505921	0.281229
Ülkü Ocakları	Government	88	21.374901	18.961487
Vatan	Opposition	413	3.565631	0.554589
Yeşil Sol Parti	Opposition	22 014	0.107690	0.000984
YRP	Government	1 637	1.468014	0.112448
Zafer	Opposition	152	8.607313	6.109121

†Not a political party but retained for completeness; bloc label marked "Neutral"

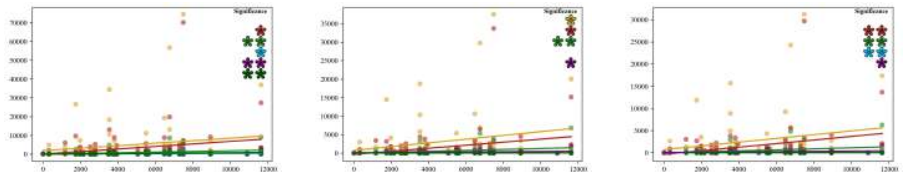

(a) Age Group, 20–25 (*raw*) (b) Age Group, 20–25 (> 0.90) (c) Age Group, 20–25 (> 0.95)

Fig. 6: Distributions for the 20–25 age group covariate, shown for raw data and with confidence filters of >0.90 and >0.95 applied.

geolocated corpus of approximately 1.22 million partisan symbols detected at a held-out accuracy of $\mathrm{mAP}_{50:95} = 0.838$. When these detections are merged with precinct-level returns, the data reveal that a marginal increase of ten thousand additional flags or banners within a district is associated with a 1.5- to 4.1-percentage-point rise in that party's vote share after controlling for month, party, district, and a comprehensive socioeconomic battery. The corresponding effect on voter turnout never exceeds two one-hundredths of a percentage point, indicating that visible symbolism reinforces partisan preferences rather than expanding the electorate.

Spatial diagnostics add further texture. Governing parties cluster their symbols more tightly than challengers: the mean convex-hull area per AKP–MHP flag is 17.4 km^2, whereas opposition parties average 23.8 km^2, a statistically decisive gap of 6.4 km^2. Formal alliances, however, do not materially affect micro-level co-location; allied and non-allied party pairs differ in mean inter-flag distance by less than a millimeter once temporal and spatial heterogeneity is absorbed. Of fourteen demographic and built-environment covariates, only the share of residents aged twenty to twenty-five exhibits a stable relationship with symbolic density: a one-standard-deviation increase in that cohort predicts an additional 0.31 symbols per panorama, whereas indicators such as educational attainment, gender balance, housing stock, and commercial density do not survive multiple-testing adjustment.

Several methodological contributions underwrite these findings. First, the pipeline scales contemporary vision transformers to panoramas as large as 16 000 × 8 000 pixels without a discernible loss of recall on objects occupying as few as forty pixels per side. Six million proposal passes and twenty-five million classifier crops were completed in eleven GPU-days—evidence that continental-scale audits are tractable on academic hardware. Second, the hand-labeled benchmark of 18,035 objects attained an inter-annotator agreement of $\kappa = 0.82$, ensuring that transfer learning rests on a reliable ground truth. Third, the entire code base, trained weights, and georeferenced outputs are released under a permissive license, removing the proprietary barrier that often inhibits replication in computational social science.

Limitations are quantifiable. Approximately 90 percent of Istanbul's precincts are represented only by imagery captured before 2016, placing an upper bound on measurement completeness and temporal resolution. Attempts to infer flag size

from monocular depth maps yielded mean absolute percentage errors between sixty-five and seventy-eight percent, rendering surface-area estimates unusable for causal inference until more accurate depth models become available.

The portability of this approach invites a series of comparative extensions. In settings characterized by a proliferation of roadside billboards, the same detector could quantify whether dominant parties cluster visual assets in their strongholds while challengers disperse theirs more widely. In jurisdictions where regulations limit the physical footprint of campaign materials, integrating symbol-density metrics with enforcement records could reveal whether compliance pressures shift resources toward digital outreach or other less-regulated forms of public-space messaging. Moreover, in contexts of mass mobilization or protest, tracking temporal changes in color saturation and symbol prevalence may serve as a real-time proxy for collective action intensity, complementing conventional measures such as crowd counts from overhead imagery.

The Istanbul results also have regulatory implications. Municipal authorities seeking a level playing field could introduce a ceiling and rely on automated audits, such as the one presented here, to issue time-stamped compliance notices rather than respond to anecdotal complaints. More generally, because the detector can be adapted to any jurisdiction with panoramic imagery, election commissions in both established and emerging democracies can incorporate visual-saturation indices into campaign-finance oversight, media-balance calculations, and studies of voter intimidation.

In sum, by converting the visual environment into structured data, this study supplies a replicable blueprint for tracking symbolic competition at scale. As panoramic coverage expands, depth estimation improves, and multi-class detectors become even more precise, the public streetscape will increasingly serve as a high-resolution barometer of partisan effort—one that scholars, journalists, and regulators can monitor in near real time across a wide spectrum of electoral regimes.

References

Adalet ve Kalkınma Partisi (AK Party) (2015). *Seçim Faaliyet Raporu 2015*. Official Party Report (in Turkish).

Ahmad, Ashfaq (2020). "Politics of Electoral Symbols in Pakistan". In: *Global Social Sciences Review* V.II, pp. 123–134. URL: https://www.humapub.com/admin/alljournals/gssr/papers/NDNBgb4DoF.pdf.

Akgun, Gokcen (2016). "Optimization of resource allocation and routing in political campaign problems". MA thesis. Wichita State University.

Akgün, Sinan Mert (2016). "Optimization of Resource Allocation and Routing in Political Campaigns: A Case Study from Turkey". In: *Wichita State University Technical Report*. Available at SOAR: http://hdl.handle.net/10057/12523.

Akyon, Fatih Cagatay, Sinan Onur Altinuc, and Alptekin Temizel (2022). *Slicing Aided Hyper Inference and Fine-Tuning for Small Object Detection*. DOI: 10.1109/ICIP46576.2022.9897990.

Atakuman, Çiğdem (2010). "Value of Heritage in Turkey: History and Politics of Turkey's World Heritage Nominations". In: *Journal of Social Archaeology* 10.3, pp. 285–313. DOI: 10.1177/1469605310379974.

Ayan, Banu (2010). "Authoritarian Party Structures in Turkey: A Comparison of the Republican People's Party and the Justice and Development Party". In: *Turkish Studies* 11.3, pp. 395–416.

Ayan, Pelin (2010). "Authoritarian Party Structures in Turkey: A Comparison of the Republican People's Party and the Justice and Development Party". In: *Turkish Studies* 11.2, pp. 197–215.

Aydın-Düzgit, Senem and Evren Balta (2017). *Turkey after the July 15 Coup Attempt: When Elites Polarize over Polarization*. Istanbul Policy Center.

Balta, Evren (2024). *Turkey's 2024 Local Elections: A Conversation with Evren Balta and Kadir Ustun*. Brandeis University Crown Center for Middle East Studies.

Baykan, Aysegul and Tali Hatuka (2010). "Politics and Culture in the Making of Public Space: Taksim Square, 1 May 1977, Istanbul". In: *Planning Perspectives* 25.1, pp. 49–68. DOI: 10.1080/02665430903421734.

Baykan, Ayşegül and Tali Hatuka (2010). "Politics and Culture in the Making of Public Space: Taksim Square, 1 May 1977, Istanbul". In: *Planning Perspectives* 25.1, pp. 49–68. DOI: 10.1080/02665430903421734.

Baykan, Toygar Sinan and Murat Somer (2022). "Politics of Notables versus National Machine: Party Organizations and Clientelism during AKP Governments". In: *European Journal of Turkish Studies* 34, online.

Bhat, Shariq Farooq, Reiner Birkl, Diana Wofk, Peter Wonka, and Matthias Müller (2023). "ZoeDepth: Zero-shot Transfer by Combining Relative and Metric Depth". In: *arXiv preprint arXiv:2302.12288*. Accessed 29 July 2025. URL: https://arxiv.org/abs/2302.12288.

Buğra, Ayşe and Ayşen Candas (2011). "Change and Continuity under an Eclectic Social Security Regime: The Case of Turkey". In: *Middle Eastern Studies* 47.3, pp. 515–528.

Cinnamon, Jonathan and Lindi Jahiu (2022). "Panoramic Street-Level Imagery in Data-Driven Urban Research: A Comprehensive Global Review". In: *Urban Studies*. Online first.

Cwalina, Wojciech, Andrzej Falkowski, and Bruce I. Newman (2011). *Political Marketing: Theoretical and Strategic Foundations*. 1st ed. New York, NY: Routledge. DOI: 10.4324/9781315701684.

Denizeau, Aurélien (2023). *Political Coalitions in Turkey in the Run-Up to the 2023 Elections*. IFRI Policy Paper.

Dorroll, Courtney (2024). *Spatial Politics in Istanbul: Turning Points in Contemporary Turkey*. Edinburgh University Press.

Edelman, Murray (1964). *The Symbolic Uses of Politics*. University of Illinois Press.

Epic Games (2023). *Panini Projection and Geometric Undistortion in Unreal Engine 5*. https://dev.epicgames.com/documentation/en-us/unreal-engine/panini-projection-in-unreal-engine. Accessed 29 July 2025.

Erensü, Sinan, Barış İne, and Yaşar Adnan Adanalı (2022). "From the Occupied Parks to the Gardens of the Nation: Politics and Aesthetics of Urban Greenery in Post-Gezi Istanbul". In: *Social Text* 40.1, pp. 109–134. DOI: 10.1215/01642472-9495146.

Esen, Berk and Şebnem Gümüşcü (2016). "Rising Competitive Authoritarianism in Turkey". In: *Third World Quarterly* 37.9, pp. 1581–1606.

Freitas, Frederico, Todd Berreth, Yi-Chun Chen, and Arnav Jhala (Aug. 2023). "Characterizing the perception of urban spaces from visual analytics of street-level imagery". In: *AI & SOCIETY* 38.4, pp. 1361–1371. ISSN: 1435-5655. DOI: 10.1007/s00146-022-01592-y. URL: https://doi.org/10.1007/s00146-022-01592-y.

Gebru, Timnit, Jonathan Krause, Yilun Wang, Jia Deng, and Li Fei-Fei (2017). "Using Deep Learning and Google Street View to Estimate the Demographic Makeup of the U.S." In: *Proceedings of the National Academy of Sciences* 114.50, pp. 13108–13113. DOI: 10.1073/pnas.1700035114.

Göksu, Oğuz (2019). "Election Strategies and Political Campaigns of the AK Party: Analysis of the Period of 2002-2014". In: *Turkish Studies-Social Sciences* 14.1, pp. 265–288.

Heper, Metin (1998). "Double-Faced State: Political Patronage and the Consolidation of Democracy in Turkey". In: *Political Parties in Turkey*. Ed. by Metin Heper and Jacob Landau. Frank Cass, pp. 102–121.

Ilic, Ljubomir, Michael Sawada, and Andrea Zarzelli (2019). "Deep Mapping Gentrification in a Large Canadian City Using Deep Learning and Google Street View". In: *PLoS ONE* 14.3, e0212814. DOI: 10.1371/journal.pone.0212814.

Kalaycıoğlu, Ersin (2022). "The Evolution of Conventional Political Participation in Turkey". In: *The Oxford Handbook of Turkish Politics*. Ed. by Güneş Murat Tezcür. Oxford University Press, pp. 715–732.

Karaveli, Halil (2017). "Authoritarian Populism and Public Rituals in Turkey". In: *Current History* 116.792, pp. 332–337.

Kutlu, Mücahid, Kareem Darwish, Cansin Bayrak, Ammar Rashed, and Tamer Elsayed (2019). "Embedding-based Qualitative Analysis of Polarization in Turkey". In: *arXiv*. eprint: 1909.10213. URL: https://arxiv.org/abs/1909.10213.

Lee, Alexander and Weihong Qi (2024). "Election symbols and vote choice: evidence from India". In: *Political Science Research and Methods* 12.4, pp. 879–887. DOI: 10.1017/psrm.2023.58.

Lees-Marshment, Jennifer (2014). *Political marketing: Principles and applications*. Routledge.

Minderer, Tobias, Alexey Dosovitskiy, Neil Houlsby, Daniel Keysers, Andreas Geiger, et al. (2022). "Simple Open-Vocabulary Object Detection with Vision Transformers". In: *arXiv preprint arXiv:2208.02536*. URL: https://arxiv.org/abs/2208.02536.

Mitchell, Don (1995). "The End of Public Space? People's Park, Definitions of the Public, and Democracy". In: *Annals of the Association of American Geographers* 85.1, pp. 108–133. DOI: 10.1111/j.1467-8306.1995.tb01797.x.

Naik, Nikhil, Scott Duke Kominers, Ramesh Raskar, Edward L. Glaeser, and César A. Hidalgo (2017). "Computer Vision Uncovers Predictors of Physical Urban Change". In: *Proceedings of the National Academy of Sciences* 114.29, pp. 7571–7576. DOI: 10.1073/pnas.1619003114.

Naik, Nikhil, Jesse Philipoom, Ramesh Raskar, and César A. Hidalgo (2014). "Streetscore—Predicting the Perceived Safety of One Million Streetscapes". In: *CVPR Workshops*, pp. 793–799.

Niffenegger, Phillip (1989). "Strategies for success from the political marketers". In: *Journal of Consumer Marketing* 6.1, pp. 45–51.

Öktem, Kerem and Karabekir Akkoyunlu (2016). "Exit from Democracy: Illiberal Governance in Turkey and Beyond". In: *Southeast European and Black Sea Studies* 16.4, pp. 469–480.

OSCE/ODIHR (2017). *Republic of Turkey — Constitutional Referendum, 16 April 2017: Final Report*. Tech. rep. Warsaw: Office for Democratic Institutions and Human Rights. URL: https://www.osce.org/odihr/elections/turkey/324816.

Özbudun, Ergun (2000). *Contemporary Turkish Politics: Challenges to Democratic Consolidation*. Lynne Rienner.

Plaggenborg, Stefan (2015). *Ordnung und Gewalt: Kemalismus–Faschismus–Sozialismus*. München: Oldenbourg.

Press, Associated (May 2021). *Erdogan opens mosque in Taksim Square, long a symbol of secular Turkey*. URL: https://apnews.com/article/middle-east-religion-government-and-politics-turkey-recep-tayyip-erdogan-4c2e3a9c971b8a2d3a7a2b0e2f2c5f5f.

Ranftl, René, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun (2020). "Towards Robust Monocular Depth Estimation: Mixing Datasets for Zero-shot Cross-dataset Transfer". In: *Proceedings of the IEEE/CVF Con-*

ference on Computer Vision and Pattern Recognition (CVPR). Accessed 29 July 2025. URL: https://arxiv.org/abs/1907.01341.

Reuters (May 2019). *Turkey orders re-run of Istanbul election won by opposition*. URL: https://www.reuters.com/article/us-turkey-election-istanbul/turkey-orders-re-run-of-istanbul-election-won-by-opposition-idUSKCN1SC1Q6.

Rozenas, Arturas and Anastasiia Vlasenko (2022). "The Real Consequences of Symbolic Politics: Breaking the Soviet Past in Ukraine". In: *Journal of Politics* 84.3, pp. 1258–1263.

Russell, Ethan (2019). *9 things to know about Google's maps data: Beyond the Map*. URL: %5Curl%7Bhttps://developers.google.com/maps/documentation/javascript/streetview%7D.

Sayarı, Sabri (2014). "Interdisciplinary Approaches to Political Clientelism and Patronage in Turkey". In: *Turkish Studies* 15.4, pp. 655–670. DOI: 10.1080/14683849.2014.985809. eprint: https://doi.org/10.1080/14683849.2014.985809. URL: https://doi.org/10.1080/14683849.2014.985809.

Sofos, Spyros A. (2021). "Space and the Emotional Topography of Populism in Turkey: The Case of *Hagia Sophia*". In: *Cogent Social Sciences* 7.1, p. 1982495. DOI: 10.1080/23311886.2021.1982495.

Somer, Murat (2015). "Crafting a National Narrative: The CHP's Vanguard Role in Turkey's Single-Party Era". In: *Authoritarianism and Resistance in Turkey*. Ed. by Esra Özyürek. Springer, pp. 33–50.

— (2019). "Turkey: The Slippery Slope from Reformist to Revolutionary Populism". In: *The Global Rise of Populism*. Stanford University Press.

— (2021). "Rivalry and Coalition-Building between Autocrats and Democrats". In: *The Oxford Handbook of Turkish Politics*. Ed. by Aylin Güney and Oğuzhan Karatekeli. Oxford University Press.

Suel, Esra, Julien Barros, Seth Flaxman, and Ricardo A. D. F. Albuquerque (2019). "Measuring Social, Environmental and Health Inequalities Using Deep Learning and Street Imagery". In: *Scientific Reports* 9, p. 6229. DOI: 10.1038/s41598-019-42036-w.

Suel, Esra, Eirik Kolding, Marco Martino, and Edward L. Glaeser (2023). "Do Poverty and Wealth Look the Same the World Over? A Comparative Study of 12 Cities Using Street Images". In: *EPJ Data Science* 12.19. DOI: 10.1140/epjds/s13688-023-00394-6.

Tokdoğan, Nagehan (2024a). "Geographies of Symbolic Power: Islamism and the Remaking of Istanbul". In: *Middle Eastern Studies* 60.1, pp. 70–97.

— (2024b). "Geographies of Symbolic Power: Islamism and the Remaking of Istanbul". In: *Middle Eastern Studies* 60.1, pp. 70–97.

Tuğal, Cihan (2013). "'Resistance Everywhere': The Gezi Park Protests in Global Perspective". In: *New Left Review* 89, pp. 123–132.

Walton, Jeremy F. (2024). "Spatial politics in Istanbul: turning points in contemporary Turkey". In: *Turkish Studies* 25.3, pp. 570–573. DOI: 10.1080/14683849.2024.2328507.

Wang, Chien-Yao, Alexey Bochkovskiy, and Hong-Yuan Mark Liao (2024). "YOLOv9: Better, Faster, and Stronger Object Detection". In: *arXiv preprint arXiv:2404.00815*. URL: https://arxiv.org/abs/2404.00815.

Wang, Fuen (2021). *Equirec2Perspec: A tool to split equirectangular panoramas into perspective views*. https://github.com/fuenwang/Equirec2Perspec. Accessed 29 July 2025.

Wilson, Andrew (2005). *Ukraine's Orange Revolution*. Yale University Press.

Wuthrich, F. Michael and Melvyn Ingleby (2020). "The Pushback Against Populism: Running on "Radical Love" in Turkey". In: *Journal of Democracy* 31.2, pp. 24–40.

Yang, Lihe, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao (2024). "Depth Anything: Unleashing the Power of Large-Scale Unlabeled Data". In: *arXiv preprint arXiv:2401.10891*. Accessed 29 July 2025. URL: https://arxiv.org/abs/2401.10891.

Yang, Lihe, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao (2024). "Depth Anything V2". In: *arXiv preprint arXiv:2406.09414*. Accessed 29 July 2025. URL: https://arxiv.org/abs/2406.09414.

Yin, Aowen, Tobias Minderer, Ben Po Liu, Koray Kavukcuoglu, and Oriol Vinyals (2023). "OWL-ViT: Open-Vocabulary Object Detection Using Vision Transformers". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 880–890. URL: https://openaccess.thecvf.com/content/CVPR2023/html/Yin_OWL-ViT_Open-Vocabulary_Object_Detection_Using_Vision_Transformers_CVPR_2023_paper.html.

A Spatial Image Sampling

Spatial subsampling is widely recommended in large-scale street-view studies to curb spatial autocorrelation and reduce the influence of repetitively photographed facades. In principle, applying such a protocol to our Istanbul corpus would be desirable. In practice, however, two severe constraints work against an aggressively conservative design.

First, the Google Street View (GSV) archive for Istanbul is highly heterogeneous: panoramas vary in camera height, heading, tilt, capture year, and—in dense neighborhoods—distance to the street centerline. These irregularities preclude a single, standardized spatial grid without introducing systematic blind spots.

Second, preliminary experiments showed that even a cautious one-in-n thinning rule would discard between 60% and 85% of available imagery, jeopardizing statistical power and the ecological validity of downstream models. We therefore retain the full harvest of approximately 748,783 panoramas for estimation and instead deploy a purpose-built robustness sample, described below, to demonstrate that the main findings are not driven by highly duplicative scenes.

Phase 1: Deduplication. We removed one image from any pair of panoramas whose camera centroids were closer than 150 m—the empirically derived threshold at which our detector can still recognize the *same* flag across views. This

radius balances the risk of double-counting with the need to preserve within-precinct heterogeneity.

Phase 2: Landmark augmentation. Uniform subsampling disproportionately underrepresents civic landmarks where political symbolism is densest. To counteract this bias, we purposively added panoramas at concentric radii of 50 m and 100 m around Hagia Sophia and the Egyptian Bazaar, capturing each location with and without digital zoom. These draws ensure that symbol-saturated environments remain visible in the evaluation set.

After filtering, 6,496 unique panoramas remained. All were manually annotated in *Roboflow* (see preceding section). This curated subset is not the analytic workhorse of the paper; rather, it serves as a stringent test bed confirming that our core results are robust to duplicate imagery, landmark bias, and the unavoidable imperfections of spatial standardization.

B Flag Size Estimation: Methods, Evaluation, and Rationale for Exclusion

Beyond simply counting partisan symbols, we initially hypothesized that the physical prominence of a banner would convey additional information about campaign resource allocation. Accordingly, we explored whether surface area could be recovered from monocular imagery. The procedure relied on depth maps produced by the model ultimately chosen in Section B.2 and on rectilinear tiles generated as described in Section B.1.

Let B denote the bounding box surrounding a detected flag and let $\hat{d}_{x,y} \in [0,1]$ be the normalized depth at pixel (x, y). We defined the tentative surface–area proxy as

$$A(B) = \sum\nolimits_{(x,y)\in B} \hat{d}_{x,y}.$$

Pixel-to-metric coefficients were calibrated on a controlled flag-validation set (Table 12).

Camera column	**Left**	**Center**	**Right**
Near	0.000188	0.000090	0.000183
Middle	0.000417	0.000818	0.000177
Far	0.002348	0.001903	0.005991

Table 12: Empirical pixel–to–metric coefficients derived from outdoor flag photographs.

In practice the approach proved untenable. Normalized monocular depth lacks a stable physical scale, so the conversion to metric units remained ad hoc. Furthermore, because depth is estimated relatively, the inferred distance of a flag shifted with scene composition: a foreground tree, for example, compressed the depth field of every other object. Finally, projection from three to two dimensions inherently discards geometric information, permitting multiple equally plausible depth configurations. Coefficients that fitted one photographic context

failed catastrophically in another, yielding surface-area estimates too noisy for inferential work. On balance, we therefore exclude flag size from the quantitative analysis and document the exercise here for transparency.

B.1 Undistortion

Rectifying equirectangular panoramas is indispensable for accurate symbol detection. We evaluated three open-source workflows on a stratified sample of 500 images: *Equirec2Perspec* (F. Wang, 2021), which applies gnomonic projection to generate perspective tiles; *Panini-based geometric undistortion* implemented in Unreal Engine 5 (Epic Games, 2023), which maps 2-D imagery onto a spherical mesh before reprojection; and *SAHI* (Akyon et al., 2022), which slices high-resolution images into overlapping tiles prior to inference. Using peak signal-to-noise ratio, edge acuity, and retention of small partisan symbols as evaluation criteria, *Equirec2Perspec* consistently out-performed the alternatives and was adopted for the production pipeline.

B.2 Depth Estimation: Manual Validation

To assess the feasibility of depth-based size recovery, we conducted a controlled validation with three American flags of known dimensions (3 ft $\times$ 5 ft, 8.3 in $\times$ 5 in, and 8 in $\times$ 12 in). Photographs were captured on an iPhone 13 Pro (26 mm, $f/1.5$, 3024 $\times$ 4032 px). The largest flag was photographed outdoors at 15, 40, and 80 m and at three elevation angles (4.57 m above eye level, eye level, and 4.57 m below). The two smaller flags were photographed indoors at 4.9, 11.3, and 19.8 m under analogous angles.

Depth Anything V1 (Yang, Kang, Huang, Xu, et al., 2024), Depth Anything V2 (Yang, Kang, Huang, Zhao, et al., 2024), MiDaS (Ranftl et al., 2020), and ZoeDepth (Bhat et al., 2023) were benchmarked on both the controlled flag photographs and a representative subset of symbol-rich street scenes. Depth Anything V2 (outdoor base) produced the lowest absolute error in object spacing and the most stable performance across lighting regimes, so we adopted it for exploratory size estimation. Two systematic errors nevertheless remained: (i) very near foreground objects compressed the depth field of the entire scene, and (ii) residual distortion in 360° tiles blurred fine details, reducing depth precision for distant flags.

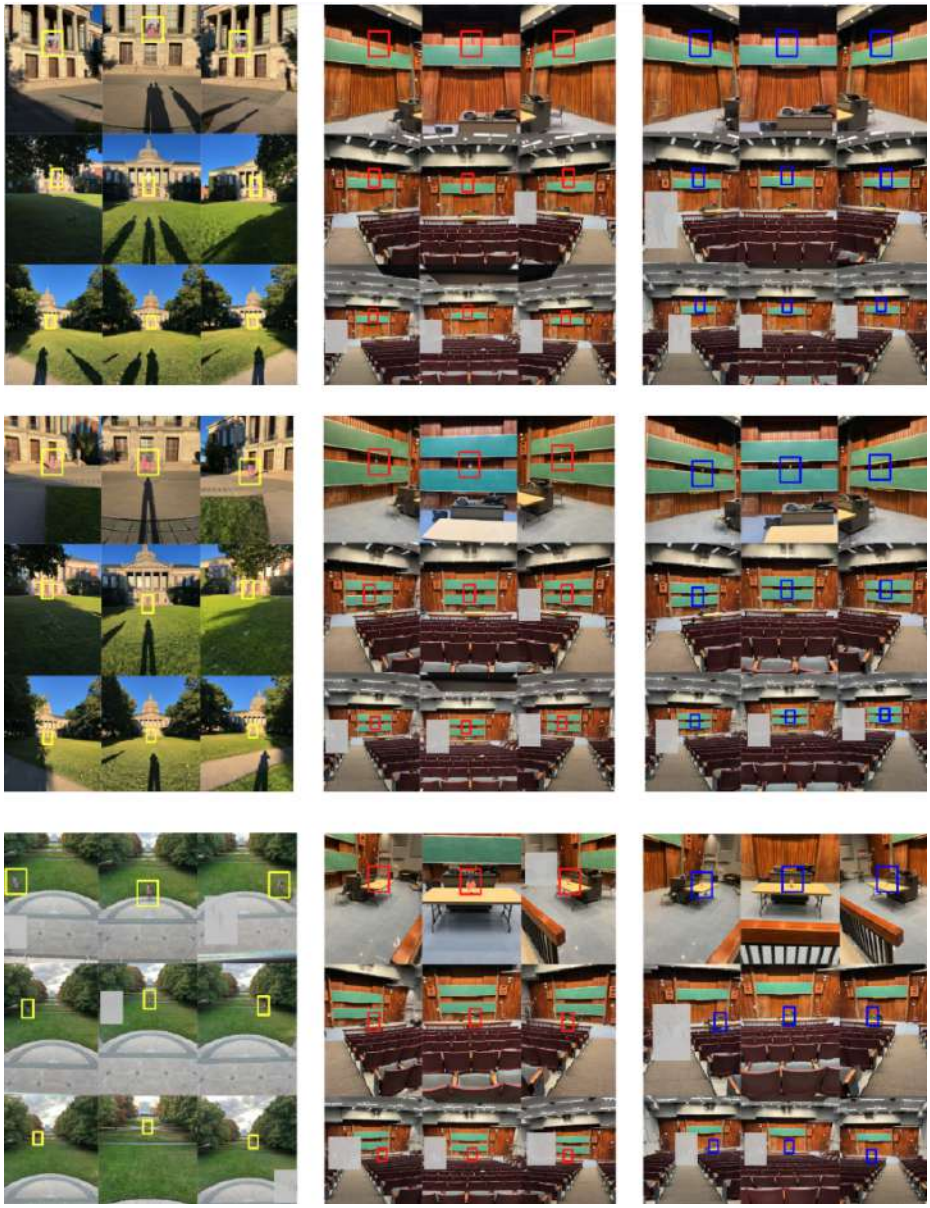

Fig. 7: Large, medium, and small flags photographed above eye level (row 1), at eye level (row 2), and below eye level (row 3).

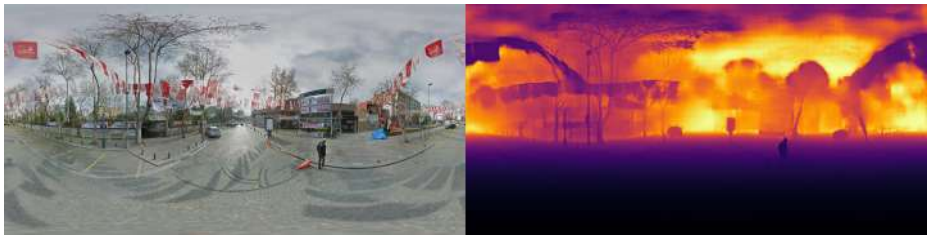

Fig. 8: Street-level scene (left) and corresponding depth map generated by Depth Anything V2 (right).

The inaccuracies associated with depth estimation are shown in Figure 9.



Fig. 9: Image 1: example of depth map where flag is too close to camera; Image 2: example of depth map where image distortion causes inaccuracies

C Overview of Street-Level Symbolic Visibility

This section presents a comprehensive overview of Google Street View panorama coverage (left subfigure), the detection of political party symbols (center subfigure), and the occurrence of Atatürk iconography and the Turkish flag (right subfigure) across all Istanbul districts.

D Pairplots for Social Variables and Party Flags



ATAŞEHİR

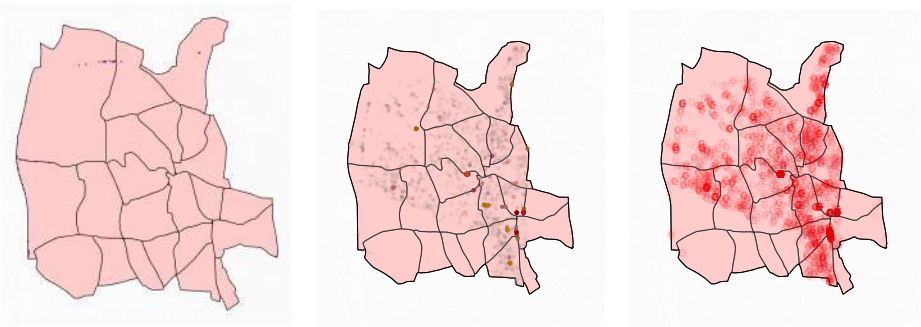

BAĞCILAR

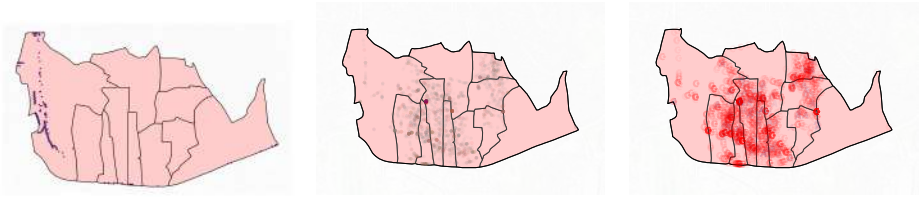

BAHÇELİEVLER

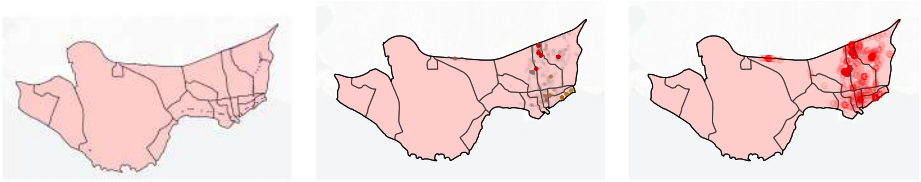

BAKIRKÖY

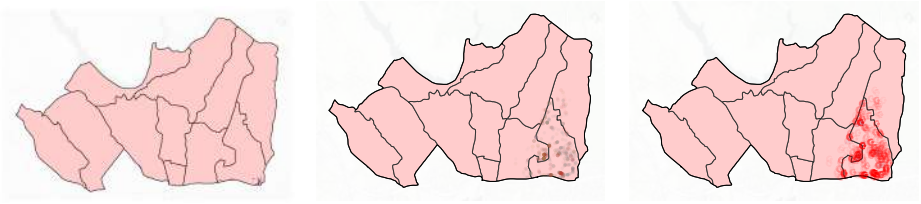

BAŞAKŞEHİR

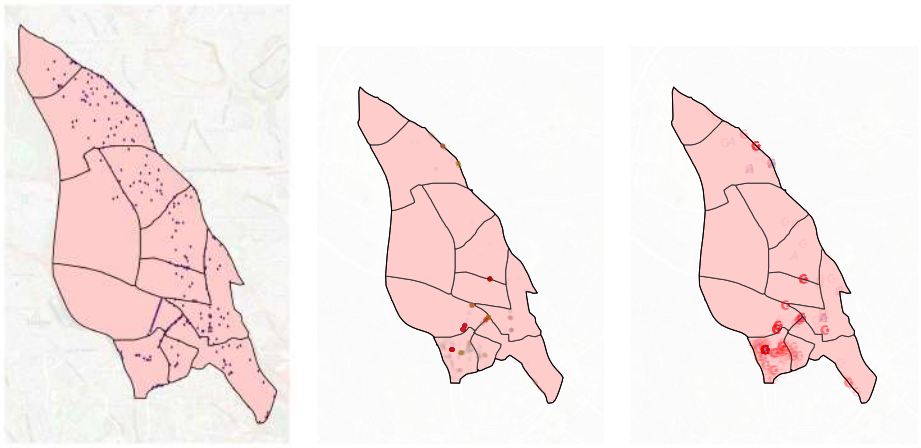

BAYRAMPAŞA

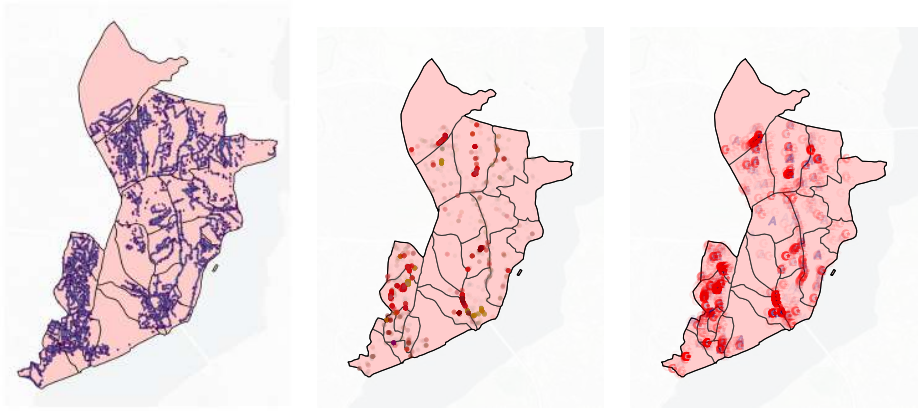

BEŞİKTAŞ

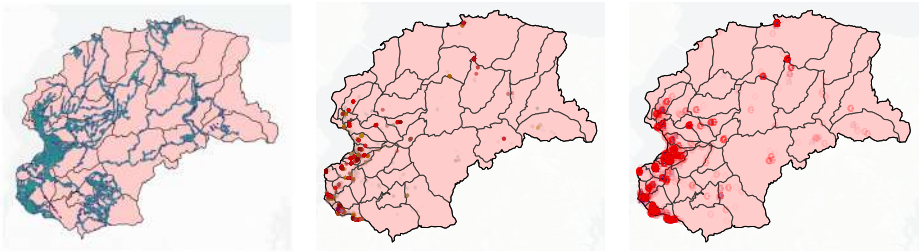

BEYKOZ

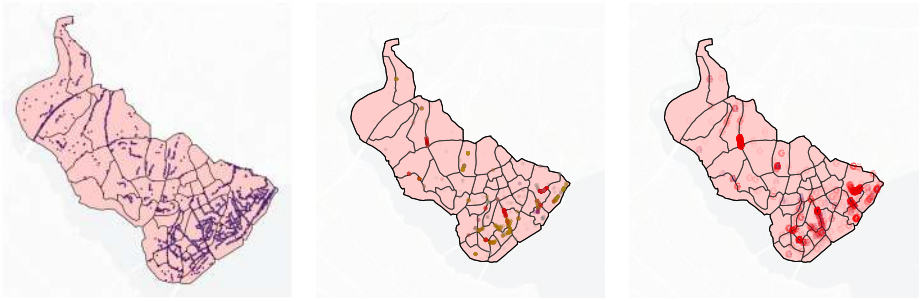

BEYOĞLU

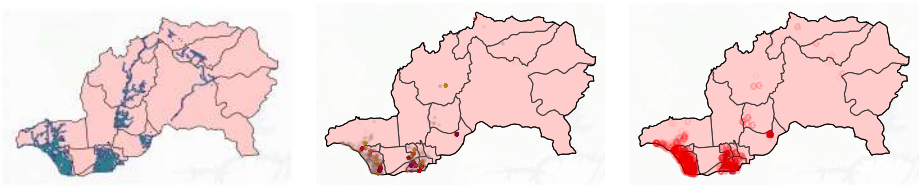

ÇEKMEKÖY

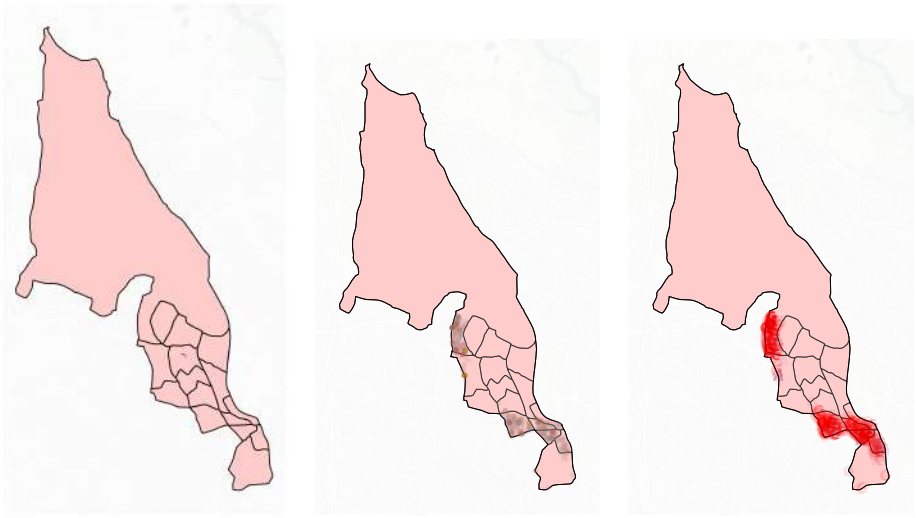

ESENLER

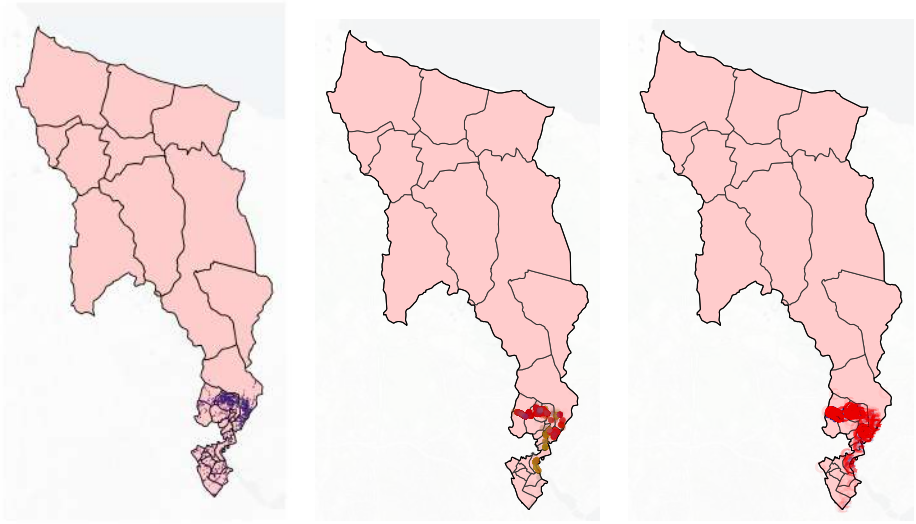

EYÜPSULTAN

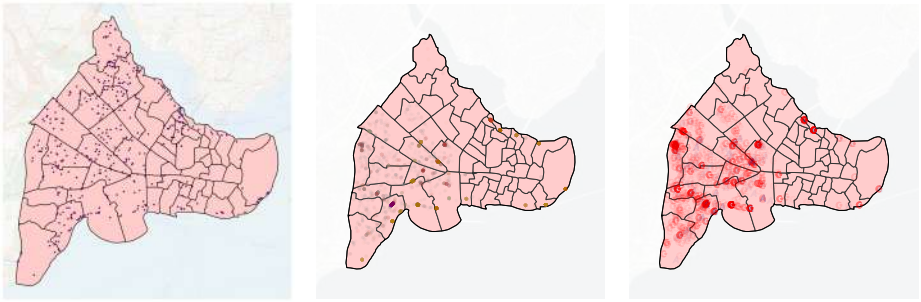

FATİH

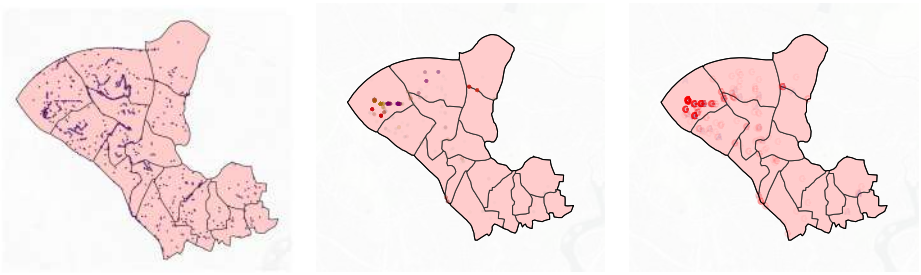

GAZİOSMANPAŞA

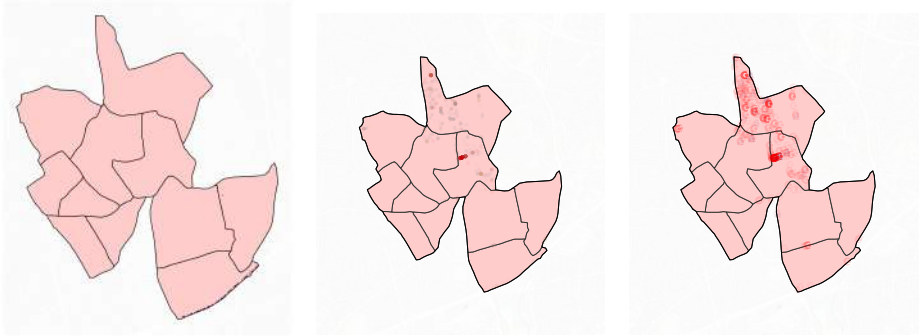

GÜNGÖREN

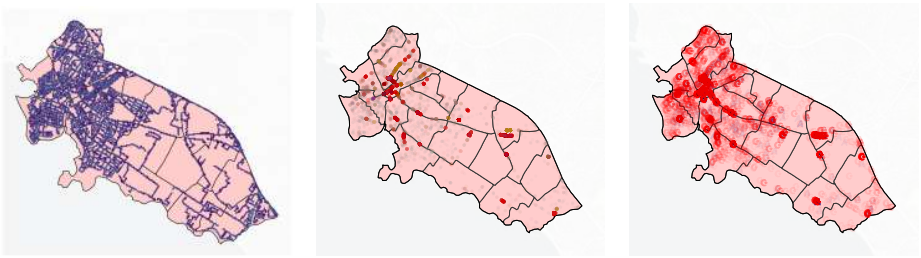

KADIKÖY

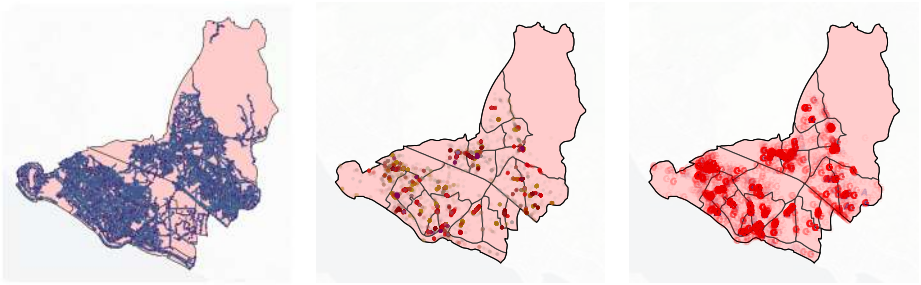

KARTAL

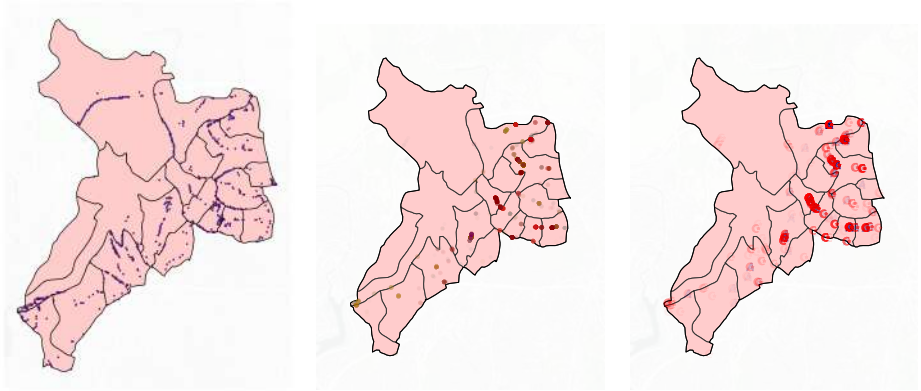

KAĞITHANE

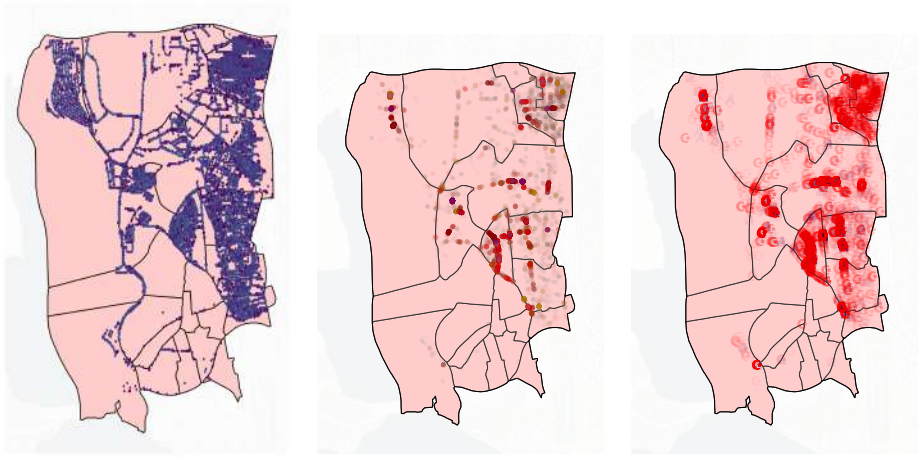

KÜÇÜKÇEKMECE

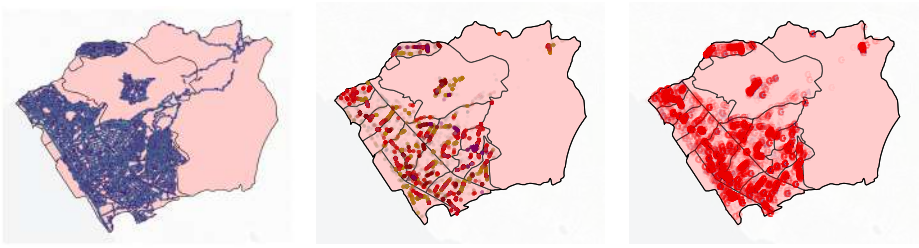

MALTEPE

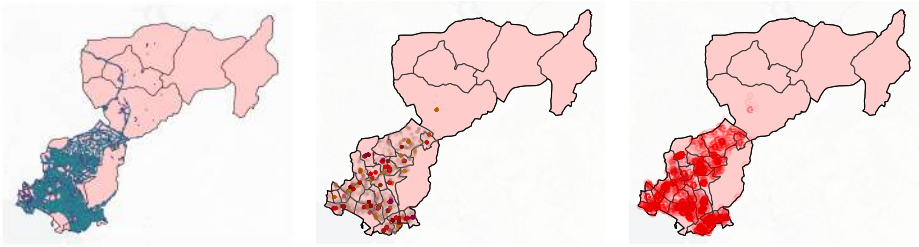

PENDİK

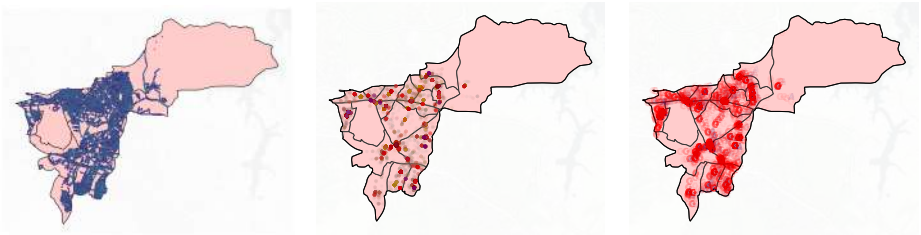

SANCAKTEPE

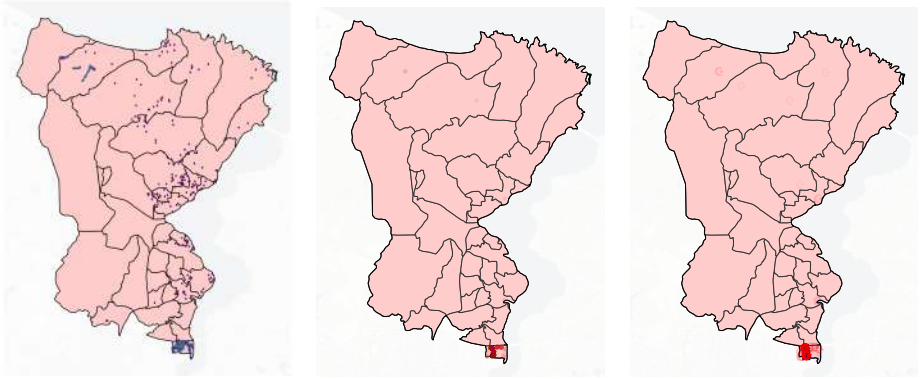

SARIYER



SULTANBEYLİ

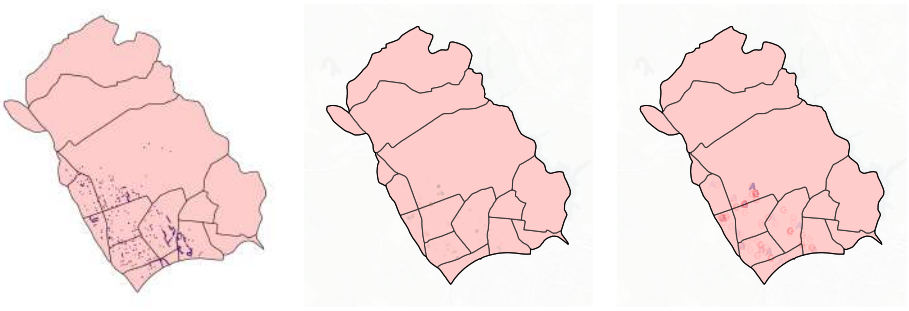

SULTANGAZİ

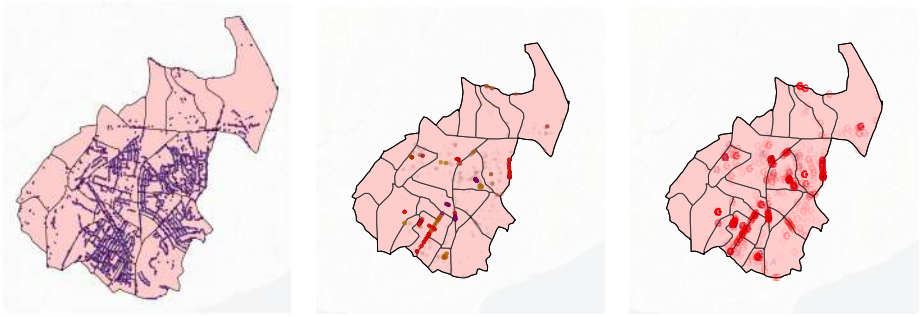

ŞİŞLİ

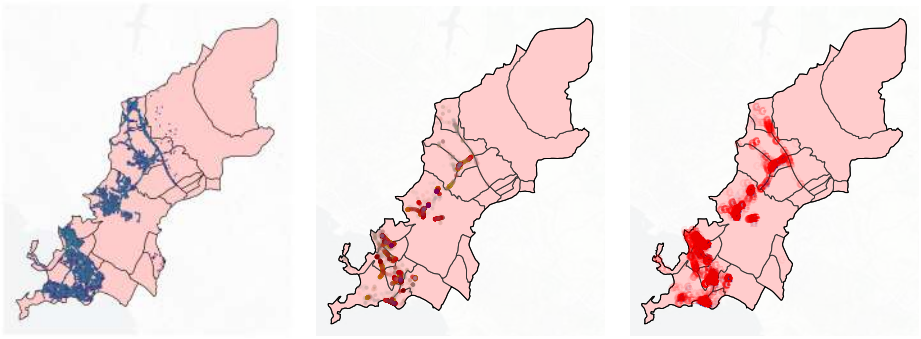

TUZLA

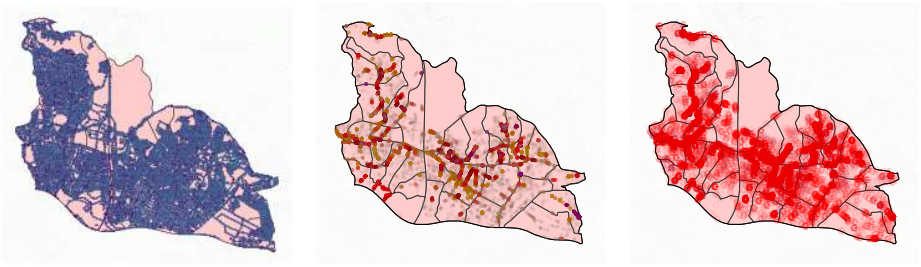

ÜMRANİYE

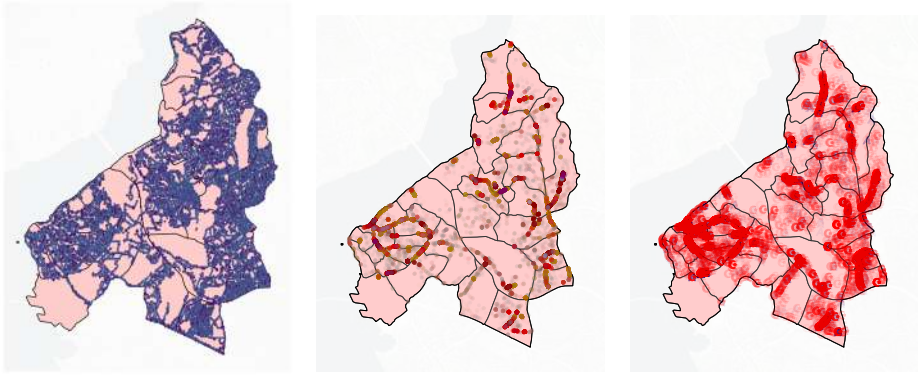

ÜSKÜDAR

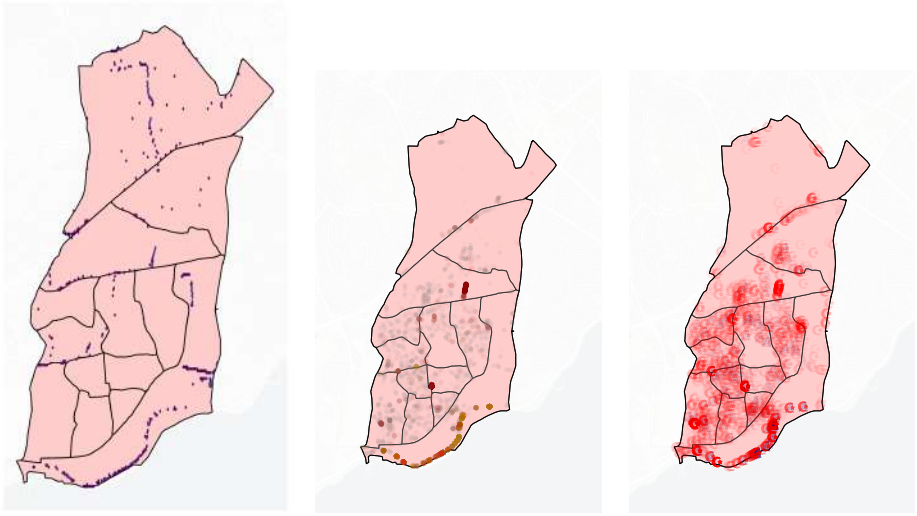

ZEYTİNBURNU

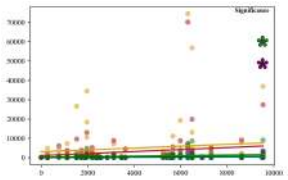

Age Group, 15–20 (*raw*)

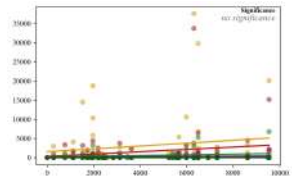

Age Group, 15–20 (> 0.90)

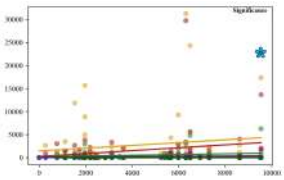

Age Group, 15–20 (> 0.95)

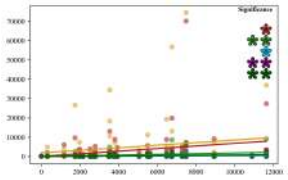

Age Group, 20–25 (*raw*)

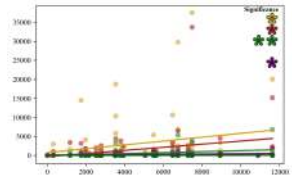

Age Group, 20–25 (> 0.90)

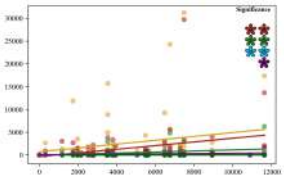

Age Group, 20–25 (> 0.95)

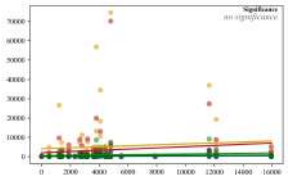

Age Group, 65+ (*raw*)

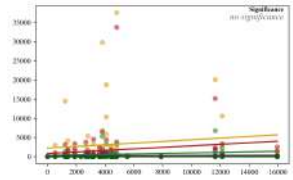

Age Group, 65+ (> 0.90)

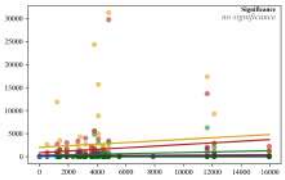

Age Group, 65+ (> 0.95)

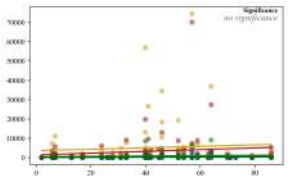

ATM (*raw*)

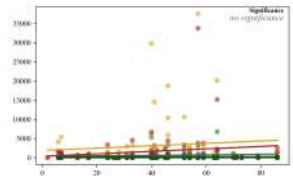

ATM (> 0.90)

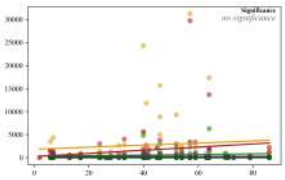

ATM (> 0.95)

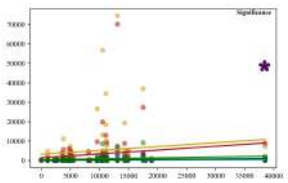

Bachelor's Degree (*raw*)

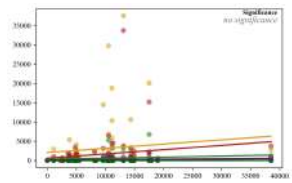

Bachelor's Degree (> 0.90)

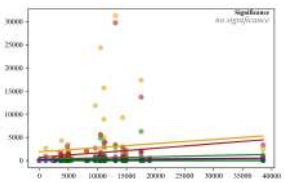

Bachelor's Degree (> 0.95)

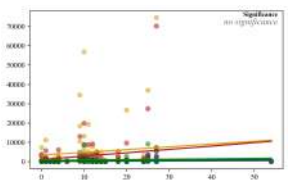

Bank (*raw*)

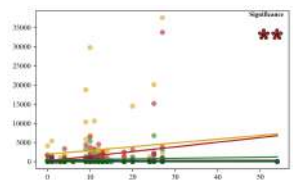

Bank (> 0.90)

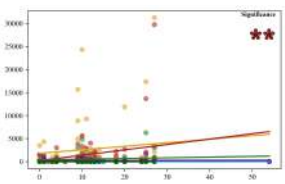

Bank (> 0.95)

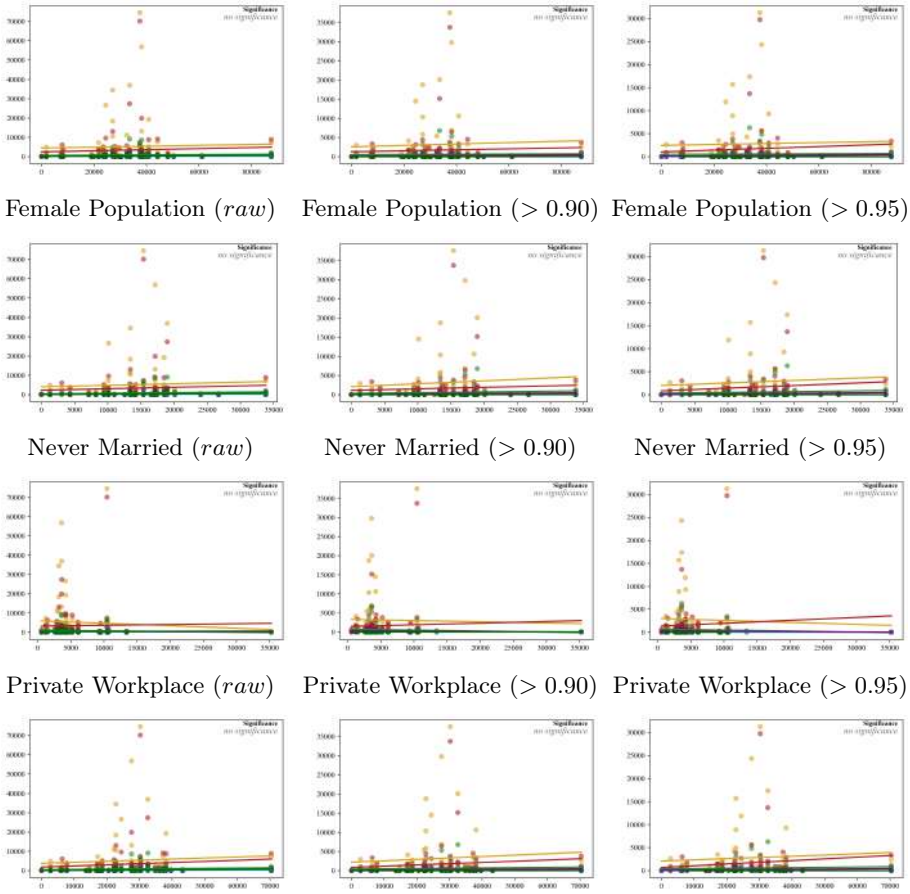

Female Population (raw) Female Population (> 0.90) Female Population (> 0.95)

Never Married (raw) Never Married (> 0.90) Never Married (> 0.95)

Private Workplace (raw) Private Workplace (> 0.90) Private Workplace (> 0.95)

Residence Count (raw) Residence Count (> 0.90) Residence Count (> 0.95)

Fig. 11: Scatterplots of party-flag density against socio-economic covariates for six major parties (AKP, CHP, İYİ, MHP, HDP, Yeşil Sol Parti). Each row is a different covariate; columns compare all detections (raw) with those exceeding confidence thresholds of 0.90 and 0.95.