Measuring AI Diffusion Across U.S. Geographies: County, State, and Metro Estimates

Amit Misra*, Jane Wang, Will Fein, Scott McCullers, Kevin White, and Juan Lavista Ferres
Microsoft AI for Good Lab
*`amitmisra@microsoft.com`

May 27, 2026

Abstract

We produce the first nationally consistent, county-level estimates of AI usage in the United States. Extending AI User Share—a population-normalized metric of active AI use—from national to subnational scales is complicated by small county samples, IP-geolocation noise, and differences in local digital infrastructure. We address these challenges by combining anonymized, aggregated Microsoft telemetry on active use of major AI services with a Bayesian spatial small-area model that borrows information across neighboring counties, state effects, and demographic covariates. County estimates are aggregated using working-age population weights to produce state, metropolitan statistical area, and urbanicity summaries. The resulting data products provide a consistent, population-normalized view of AI adoption across U.S. geographies.

Keywords: AI adoption, county-level estimation, digital divide, Bayesian spatial modeling, small area estimation

1 Introduction

National-level estimates show that AI adoption in the United States is widespread, with roughly 31% of the working-age population actively using AI tools [1]. But national averages conceal substantial geographic variation. AI usage is shaped by local digital infrastructure, demographics, industry composition, and the presence of universities and knowledge-intensive employers—none of which are distributed evenly across the country.

No existing public dataset that we are aware of provides consistent AI adoption estimates at the U.S. county level. County-level measurement is complicated by small samples in low-population areas, IP-geolocation noise, tourism, and differences in local digital infrastructure. Raw county-level telemetry signals are too noisy to use directly.

To address this gap, we extend AI User Share—the proportion of a relevant population actively using AI tools—from the national scale [1] to U.S. counties and county-equivalent areas. We combine Microsoft telemetry on active AI usage with a Bayesian spatial small-area model that borrows information from neighboring counties, statewide patterns, age composition, and urbanicity, then apply device-access adjustments, infrastructure scaling, and national calibration. County estimates are aggregated with working-age population weights to produce state, MSA, and urbanicity summaries.

2 Methods

2.1 Data and Measurement

AI User Share. Our outcome variable is **AI User Share**—the estimated proportion of a geography's working-age population that actively uses AI-powered tools. The metric was originally developed for cross-country comparisons [1]. At the county level, the raw telemetry signal is the share of stable Microsoft users engaging with major AI services; the final estimate further incorporates device-access scaling, technology-access scaling, and national calibration. For the released U.S. outputs, county estimates are re-normalized and aggregated using ACS 2024 ages 15–64 working-age population weights [2].

Microsoft telemetry. Our primary data source is anonymized, aggregated Microsoft telemetry, predominantly from desktop/laptop platform usage, covering engagement with major AI services including ChatGPT, Google Gemini, Anthropic Claude, Microsoft Copilot, and others. We restrict to "stable" users—those with at least 90 minutes of usage time in a given month—to filter transient or bot traffic. County attribution uses IP geolocation, which introduces spatial noise near county boundaries that our model accounts for.

Demographic, infrastructure, and geography data. We supplement telemetry with:

- **Age composition**: Age distributions (shares aged >65, <17, and 18–24) from the American Community Survey [3], and ACS 2024 B01001 ages 15–64 for release weighting [2].
- **Urbanicity**: Metropolitan, Micropolitan, and Rural classifications derived from the USDA Rural-Urban Continuum Codes (RUCC), which build on OMB metropolitan and nonmetropolitan concepts [4]. RUCC values of 1–3 are assigned Metropolitan, 4–7 Micropolitan, and 8–9 Rural. Metropolitan corresponds to counties in metropolitan areas; the Micropolitan label is shorthand for nonmetro counties with urban populations from 5,000 to 49,999; Rural covers counties that are completely rural or have urban populations below 5,000.
- **Technology infrastructure**: Microsoft broadband usage rates [5] and Census computer-and-internet access [3], used to adjust for differences in digital infrastructure across counties.
- **Aggregation geographies**: Census/OMB CBSA delineation files for MSAs [6] and RUCC-derived county classes for urbanicity summaries.

2.2 Modeling Approach

The fundamental challenge of county-level estimation is that many counties have too little data to produce reliable estimates on their own. Our model addresses this through adaptive shrinkage: for counties with abundant data, the estimate closely tracks what we observe; for counties with sparse data, the model relies more heavily on what neighboring counties, age composition, and urbanicity suggest the rate should be.

How the model works. Conceptually, our model estimates each county's AI adoption rate as a blend of four sources of information:

1. **The county's own data**—what we directly observe in Microsoft telemetry.
2. **Neighboring counties**—the spatial component, which assumes that nearby counties tend to have similar adoption rates.
3. **Statewide patterns**—a state-level effect that captures policy, cultural, and economic factors shared across a state.
4. **Age composition and urbanicity predictors**—local age structure and county type, which are systematically related to AI adoption.

The model automatically learns how much weight to give each source. Large, well-measured counties like Maricopa, AZ or New York County, NY are largely determined by their own data. Small rural counties like Harding, NM (population 635) are determined almost entirely by their neighbors, age composition, and urbanicity.

Formal specification. For county i, the model is:

$$\begin{aligned} y_i &\sim \text{Beta}(\mu_i \kappa_i, (1-\mu_i)\kappa_i)\,, \\ \text{logit}(\mu_i) &= \underbrace{\alpha}_{\text{national baseline}} + \underbrace{\mathbf{x}_i^\top \boldsymbol{\beta}}_{\text{age and urbanicity}} + \underbrace{s_{\text{state}(i)}}_{\text{state effect}} + \underbrace{u_i^*}_{\text{spatial}} + \underbrace{\varepsilon_i}_{\text{noise}}\,. \end{aligned} \tag{1}$$

Here y_i is the county-level telemetry signal derived from stable Microsoft users' AI engagement, μ_i is the latent adoption rate we want to estimate, and κ_i is a precision parameter that scales with sample size—giving the model more confidence in counties with more data. The covariates $\mathbf{x}_i$ include population shares by age group and urbanicity indicators. The spatial component u_i^* uses a Gaussian Markov random field with Queen contiguity on county boundaries (∼18,000 edges), an intrinsic conditional autoregressive (ICAR) prior that encourages nearby counties to have similar estimates [7].

Inference. We estimate the model using the No-U-Turn Sampler (NUTS) [8], a modern Markov chain Monte Carlo algorithm, implemented in NumPyro [9] with GPU acceleration. This produces 4,000 posterior draws (4 chains × 1,000 each) from which we derive both point estimates and uncertainty intervals for every county. For the fitted county rates, diagnostic checks showed tail effective sample sizes above 1,000, $\hat{R}$ below 1.01, and no divergent transitions.

2.3 Post-Processing Pipeline

The Bayesian model produces county-level relative AI adoption estimates. Additional post-processing steps transform these into the final county AI User Share:

1. **Replacement of unreliable estimates.** Counties can be flagged where geolocation noise, outlier cities, or anomalously high telemetry coverage relative to resident population compromise the spatial prediction. For these counties, the pipeline uses a simpler age-composition/state estimate. In the current release, 268 of 3,143 county units (8.5%) are replaced; these are predominantly smaller-population counties (median ∼12,000 working-age residents).
2. **Device and technology access scaling.** The telemetry used for the county signal is primarily desktop/laptop platform usage, but AI usage also occurs on mobile devices. We adjust final county estimates for local desktop/laptop access, smartphone-to-desktop/laptop ownership ratios from ACS S2801 [10], broadband usage [5], and Census computer-and-internet access [3].
3. **National calibration.** The model identifies relative differences between counties; calibration sets the overall level. We apply a multiplicative re-normalization so the ACS 15–64 working-age-weighted national mean equals the national AI User Share estimate.
4. **Aggregation to release geographies.** County estimates are aggregated to states, metropolitan statistical areas (MSAs), and urbanicity groups (Metropolitan, Micropolitan, Rural) as working-age-population-weighted means, using ACS 2024 ages 15–64 weights [2]. MSAs use the Census/OMB July 2023 CBSA delineation [6]; urbanicity groups follow USDA RUCC classes [4].

2.4 Validation

We validated our estimates through two independent checks. No county-level ground truth exists for AI adoption, so our validation strategy tests internal consistency, generalizability, and robustness to modeling choices.

Leave-one-state-out cross-validation. We refit the model 51 times, each time completely removing one state and predicting its counties using only nationally learned demographic relationships—no spatial information or state-level effects from the held-out state. This is a stringent test of whether the model generalizes to entirely unseen geography.

Across 1,602 large held-out counties (those with >10,000 stable users), we push both LOSO and full-data predictions through the complete post-processing pipeline so that comparisons reflect the final AI User Share. The LOSO predictions correlate at $r = 0.88$ with the full-model estimates on final outputs (MAE = 2.0 pp), indicating that age composition, urbanicity, and nationally learned relationships recover most of the county-level variation even without in-state data. The median per-state correlation is 0.93 across the 43 states with at least 10 large counties.

Spatial-block validation. LOSO validates covariate generalization but not the ICAR spatial component. To test whether spatial smoothing adds genuine predictive value, we hold out contiguous blocks of counties (∼20% of each state's counties) within each of the five largest states by population and predict using only covariates and state random effects. The spatial model's RMSE beats the non-spatial baseline in 3/5 states (CA, FL, PA; improvement 10–18%); the non-spatial baseline performs comparably or slightly better in TX and NY.

2.5 Industry Analysis

The industry results are estimated separately from the Bayesian small-area model. We fit an ordinary least squares regression of final county AI User Share on county industry employment shares from the BLS Quarterly Census of Employment and Wages [11]. Candidate sectors were modeled one at a time with regression controls for RUCC urbanicity, age composition, and log median household income from the ACS [3]. Figure 6 reports the resulting industry coefficients with 95% confidence intervals.

3 Results

The state-level map in Figure 1 shows clear patterns. AI usage is highest in Maryland, Utah, Texas, Virginia, and New Jersey. States with lower levels of AI usage include West Virginia, Maine, Montana, Mississippi, and Vermont. These states tend to have smaller metropolitan footprints, older populations, and larger rural or small-town shares.

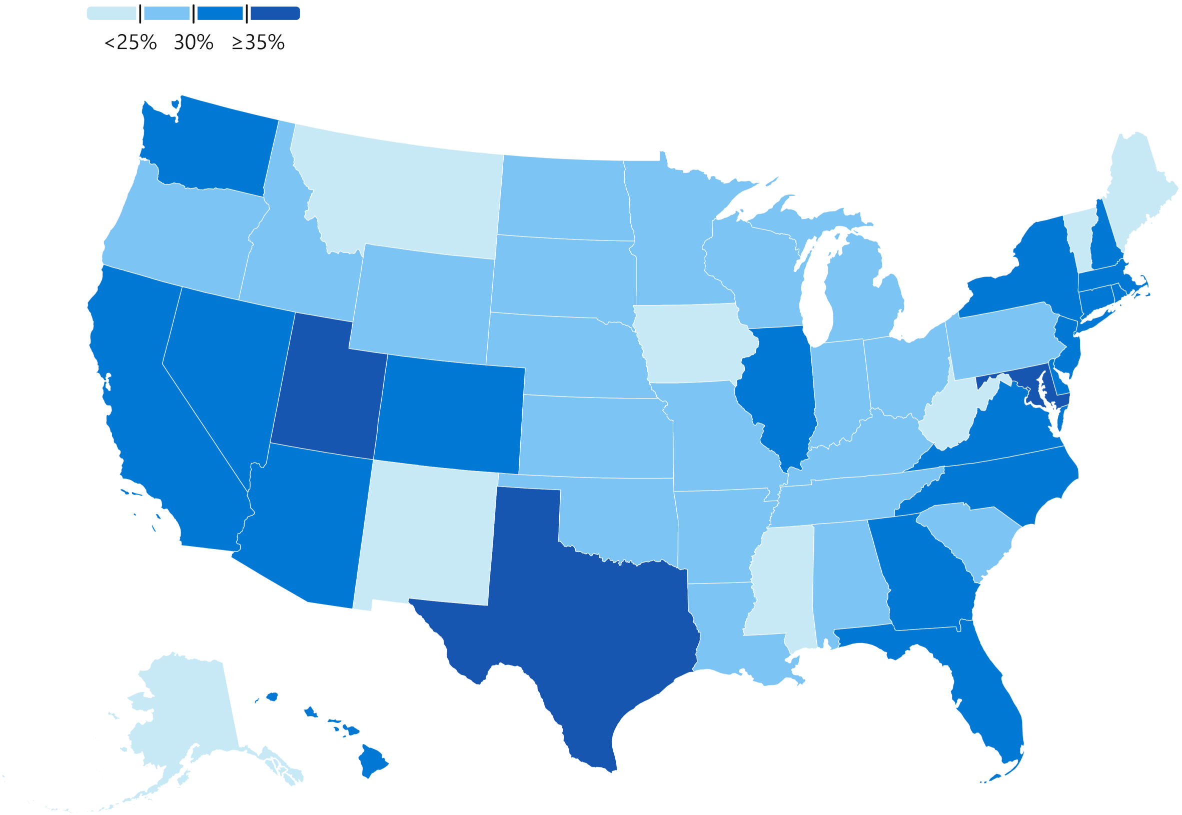

Figure 1: AI User Share by state. Darker blue indicates higher estimated AI usage, highlighting higher adoption in several Mid-Atlantic, Mountain West, and Sun Belt states and lower adoption across more rural or older-population states.

County-level estimates, shown in Figure 2, allow analysis of relationships between AI adoption and local demographic, economic, and infrastructure characteristics at a granularity that state-level data cannot resolve. A college town, a rural county, and a major metropolitan core can sit within the same state and look nothing alike; state averages blend them together. For example, at a high level, the county-level map shows that urban areas tend to have higher AI User Share. Counties with college towns also stand out. Among the top 15 counties by AI User Share with at least 10,000 working-age residents, shown in Figure 3, 13 are in the 95th percentile or higher for 18–24-year-old population share.

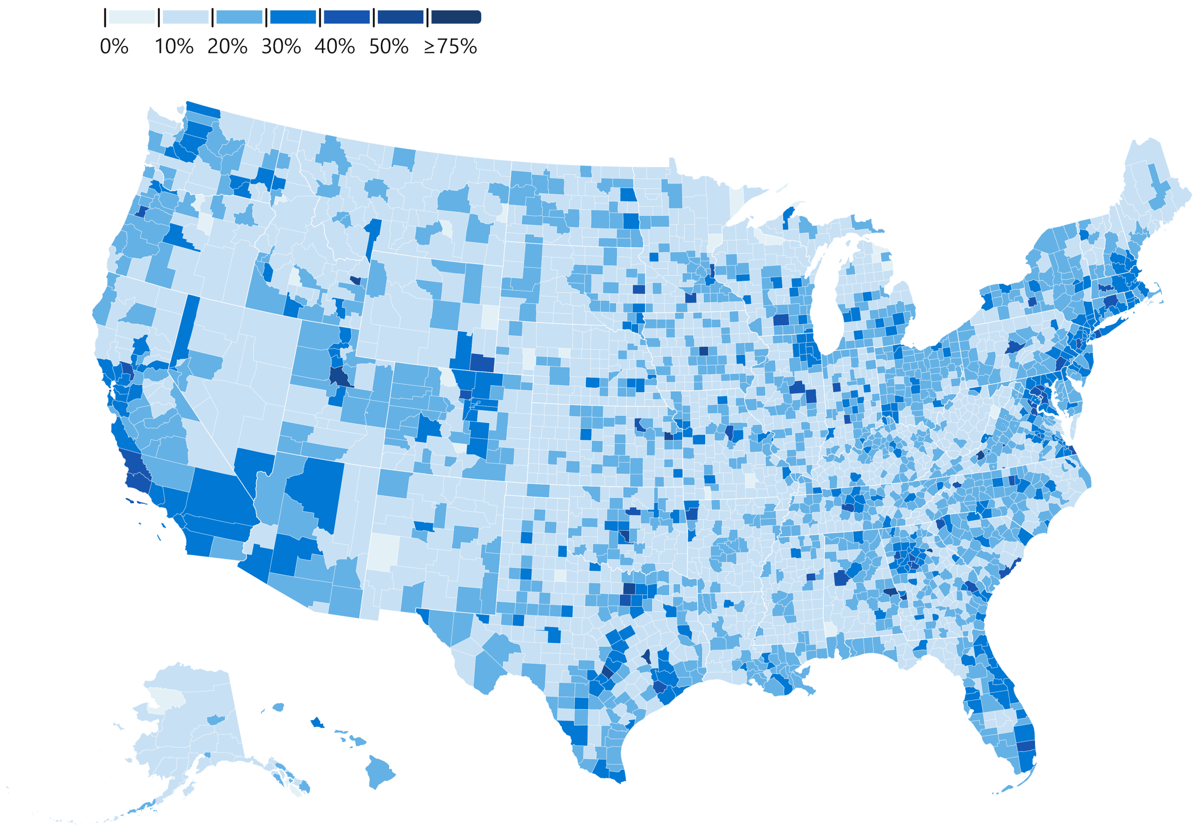

Figure 2: County-level AI User Share across the U.S.; darker blue indicates higher estimated usage. The map shows substantial within-state variation that is hidden in state averages.

Rank	County-Equivalent Area	State	Age 18–24	AI User Share	Anchor Institution
1	Williamsburg	Virginia	36.2%	73.2%	College of William & Mary
2	Harrisonburg	Virginia	32.7%	67.5%	James Madison University
3	Madison	Idaho	46.4%	67.2%	Brigham Young University–Idaho
4	Brazos	Texas	26.4%	64.0%	Texas A&M University
5	Story	Iowa	30.6%	63.8%	Iowa State University
6	Denton	Texas	8.9%	60.1%	University of North Texas
7	Clarke	Georgia	26.0%	59.0%	University of Georgia
8	Douglas	Kansas	23.0%	57.2%	University of Kansas
9	Charlottesville	Virginia	20.1%	51.7%	University of Virginia
10	Cleveland	Oklahoma	15.0%	51.6%	University of Oklahoma
11	Hays	Texas	14.2%	51.4%	Texas State University
12	Lee	Alabama	16.5%	50.8%	Auburn University
13	Montgomery	Virginia	27.8%	50.8%	Virginia Tech
14	Utah	Utah	16.8%	50.2%	Brigham Young University
15	Fredericksburg	Virginia	18.0%	49.3%	University of Mary Washington

Figure 3: Top 15 county-equivalent areas by AI User Share among areas with working-age population of at least 10,000. The table also shows 18–24 population share and nearby anchor colleges or universities, which are common among the highest-adoption counties.

One key finding from this work is a stark urban-rural divide. As shown in Figure 4, Metropolitan AI User Share is more than twice the rural share, suggesting deep differences in adoption beyond what national averages capture.

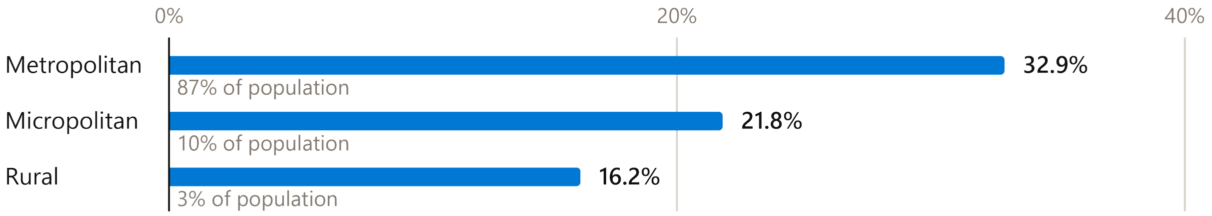

Figure 4: AI User Share by urban class. Metropolitan counties have the highest share (32.9%), followed by Micropolitan (21.8%) and Rural counties (16.2%); labels also show each class's share of the working-age population.

This pattern is especially visible in college towns. Figure 5 shows a strong correlation between 18–24 population share and AI User Share, with the highest 18–24 share counties having AI User Shares more than three times those of the lowest 18–24 share counties.

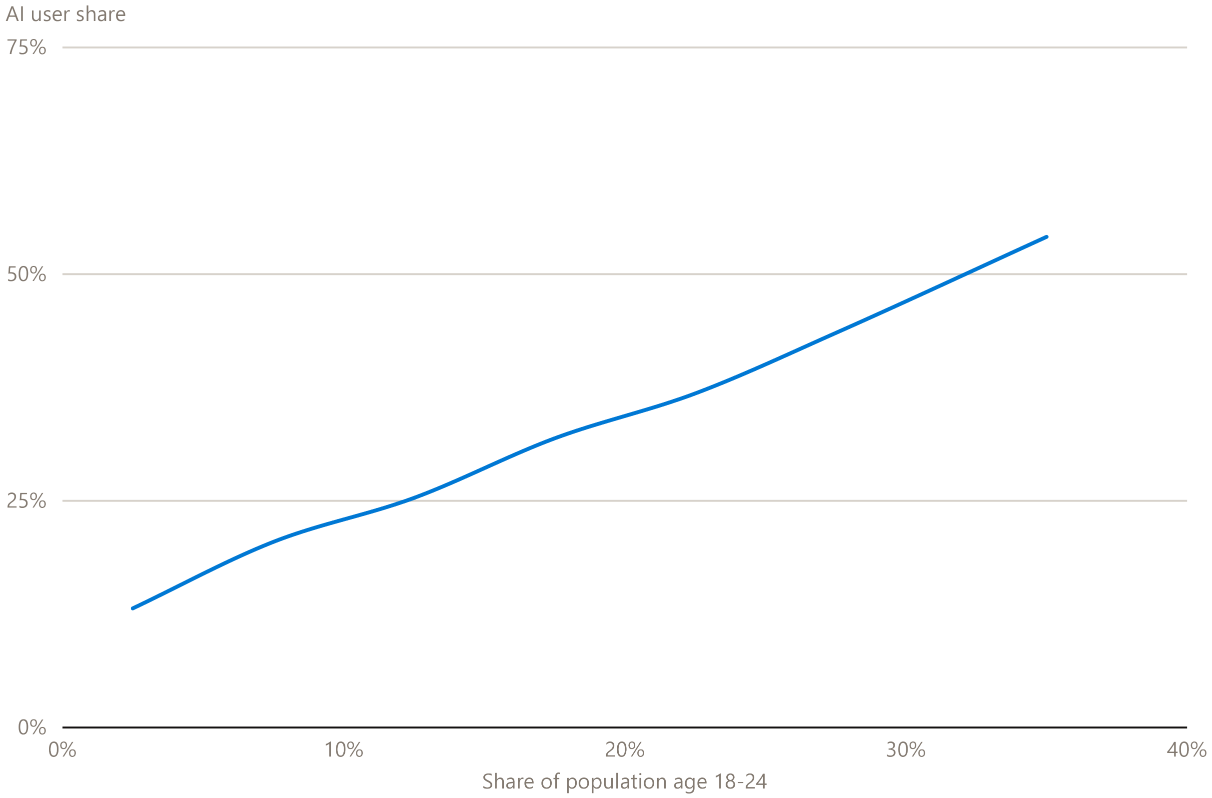

Figure 5: Mean county AI User Share by 18–24 population-share bin. The line shows a steady increase in estimated AI usage as the young-adult share rises.

Industry mix is also correlated with AI User Share, as seen in Figure 6. Counties with larger shares of professional and technical services, information work, health care, and finance show higher AI User Share; counties with larger shares of manufacturing, agriculture, construction, and mining show lower AI User Share. These associations hold in the separate OLS analysis after accounting for urbanicity, age composition, and income.

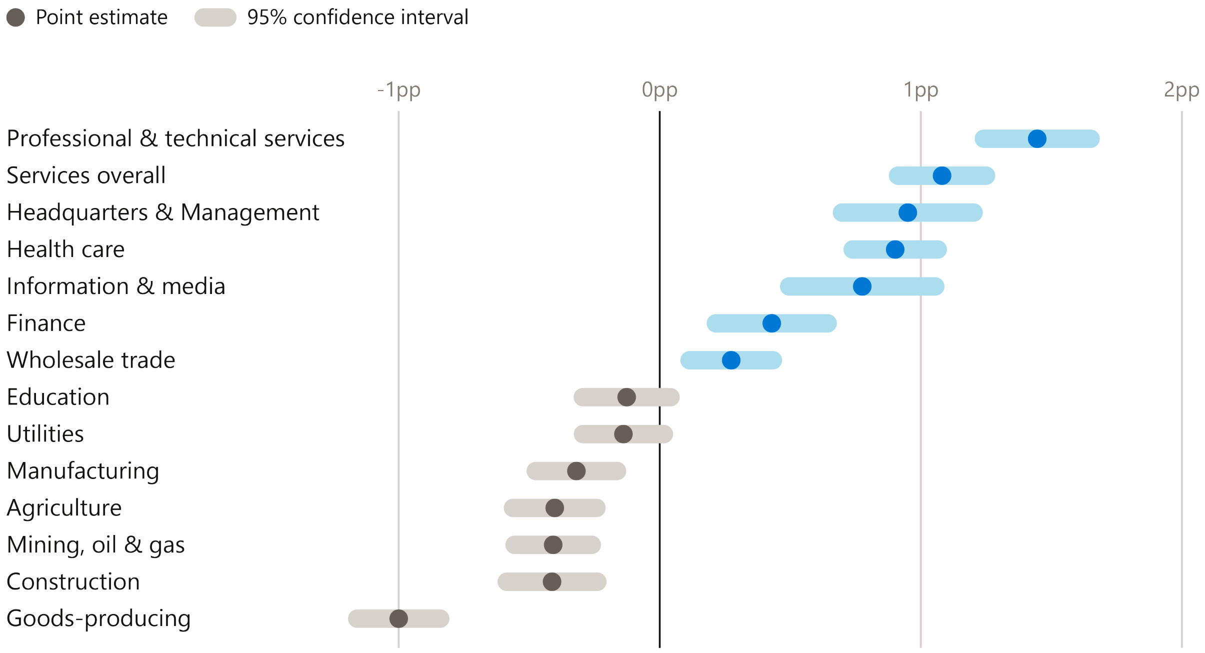

Figure 6: OLS coefficients for industry-category associations with AI User Share, controlling for urban class, age composition, and median household income. Each industry is estimated in a separate regression. Points show estimates and bars show 95% confidence intervals. Services overall are positively correlated with AI User Share, while goods-producing industries are negatively correlated.

The county-level estimates can also be aggregated to metropolitan statistical areas (MSAs) to reveal patterns within and across urban centers. Figure 7 shows the top 35 MSAs by population. The ranking shows that large population size does not explain AI adoption. High usage is also not limited to traditional technology hubs. The top-ranked metros illustrate this: Washington (#1) likely reflects the concentration of federal, legal, consulting, and professional work. Dallas and Houston point to stronger adoption in large Texas metros with corporate and energy-sector employment. Miami and Atlanta show higher usage in fast-growing Sun Belt metros with large service and business workforces.

Many of the lower-ranked metros are older industrial or slower-growth regions with strong institutions and manufacturing bases. Their lower usage rates suggest that these strengths do not automatically produce broad AI adoption. The difference appears to reflect the mix of jobs, age composition, and urbanicity that are associated with higher AI usage, as described above.

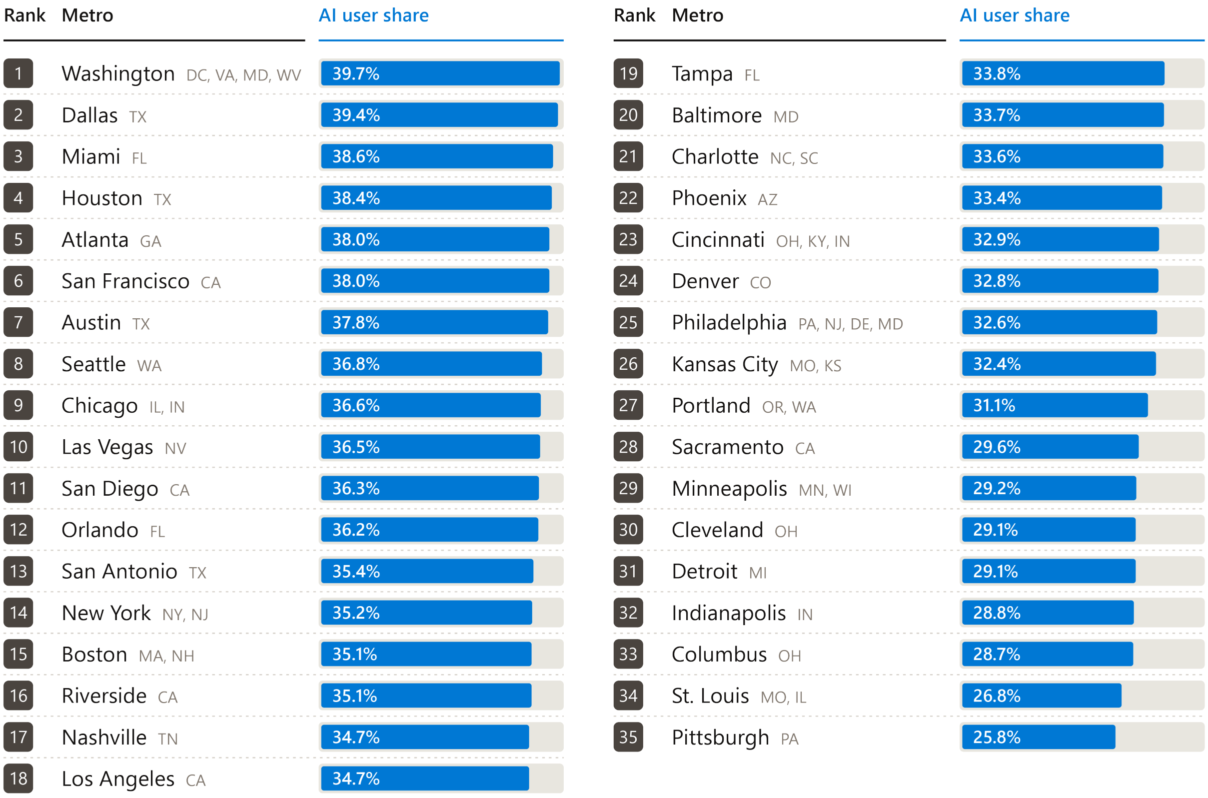

Rank	Metro	AI user share
1	Washington DC, VA, MD, WV	39.7%
2	Dallas TX	39.4%
3	Miami FL	38.6%
4	Houston TX	38.4%
5	Atlanta GA	38.0%
6	San Francisco CA	38.0%
7	Austin TX	37.8%
8	Seattle WA	36.8%
9	Chicago IL, IN	36.6%
10	Las Vegas NV	36.5%
11	San Diego CA	36.3%
12	Orlando FL	36.2%
13	San Antonio TX	35.4%
14	New York NY, NJ	35.2%
15	Boston MA, NH	35.1%
16	Riverside CA	35.1%
17	Nashville TN	34.7%
18	Los Angeles CA	34.7%
19	Tampa FL	33.8%
20	Baltimore MD	33.7%
21	Charlotte NC, SC	33.6%
22	Phoenix AZ	33.4%
23	Cincinnati OH, KY, IN	32.9%
24	Denver CO	32.8%
25	Philadelphia PA, NJ, DE, MD	32.6%
26	Kansas City MO, KS	32.4%
27	Portland OR, WA	31.1%
28	Sacramento CA	29.6%
29	Minneapolis MN, WI	29.2%
30	Cleveland OH	29.1%
31	Detroit MI	29.1%
32	Indianapolis IN	28.8%
33	Columbus OH	28.7%
34	St. Louis MO, IL	26.8%
35	Pittsburgh PA	25.8%

Figure 7: AI User Share across the 35 most populous metropolitan statistical areas (MSAs), aggregated from county estimates using working-age population weights. High usage is not limited to traditional technology hubs, and metro population size alone does not explain adoption.

4 Discussion and Conclusion

The county-level patterns align with independent survey evidence. The steep gradient in AI User Share with 18–24 population share is consistent with survey work from Pew and Gallup showing higher AI and ChatGPT use among younger adults [12, 13]. The urban-rural gap echoes survey evidence on differential AI trust and knowledge across U.S. communities [14], suggesting that the access and demographic factors captured in our model may be reinforced by attitudinal differences our model does not directly measure. The industry coefficients—positive for professional and technical services, information, finance, and health care; negative for goods-producing sectors—are also consistent with NBER evidence that workplace generative AI adoption is concentrated in professional, analytical, and other knowledge-intensive work [15]. Together, these external alignments increase confidence that the county estimates capture real variation rather than artifacts of the telemetry pipeline.

Even with this external support, three limitations are worth flagging. First, all estimates derive from anonymized, aggregated Microsoft telemetry, and the Microsoft user population may differ from the general population due to product usage, device access, and opt-out behavior; the mobile and technology access adjustments partially compensate but cannot fully remove platform-specific selection bias. Second, no independent county-level AI adoption data exists for external validation, so our internal checks and the survey alignments above test consistency and plausibility but cannot rule out biases shared across all counties. Third, several pipeline steps involve judgment, including sample-size adjustments, replacement thresholds, and national calibration; while sensitivity checks show aggregate products are stable, local county anomalies can remain. The estimates should also be read as a measurement layer across U.S. geographies, not as causal claims: associations with young-adult population share, urbanicity, or industry mix are patterns in the data that can motivate questions but do not by themselves identify why adoption differs across places.

Despite these caveats, the granularity of these estimates makes visible the patterns that national or state averages cannot resolve. The urban-rural divide, the college-town effect, and the role of industry mix all emerge clearly at the county level. They also highlight both the uneven distribution of current AI adoption and the potential opportunities to broaden it. Policymakers cannot change the urban-rural classification or industry mix of a county overnight, but they can influence potential drivers of lower AI adoption like digital infrastructure, AI literacy, and workforce training.

Data availability. Aggregated county, state, and MSA-level estimates are available at `https://github.com/microsoft/ai-diffusion-report/tree/main/data/US`.

References

[1] A. Misra, J. Wang, S. McCullers, K. White, and J. Lavista Ferres, "Measuring AI diffusion: A population-normalized metric for tracking global AI usage," *arXiv:2511.02781*, November 2025. doi: 10.48550/arXiv.2511.02781; `https://arxiv.org/abs/2511.02781`.

[2] U.S. Census Bureau, "B01001: Sex by age," 2024. 2020–2024 American Community Survey 5-Year Estimates, Detailed Tables; `https://data.census.gov/table/ACSDT5Y2024.B01001`.

[3] U.S. Census Bureau, "American community survey 5-year data (2009–2024)," January 2026. Census.gov; `https://www.census.gov/data/developers/data-sets/acs-5year.html`.

[4] U.S. Department of Agriculture, Economic Research Service, "2023 rural-urban continuum codes," 2023. Updated Jan. 22, 2024; `https://www.ers.usda.gov/data-products/rural-urban-continuum-codes`.

[5] Microsoft, "United states broadband usage percentages dataset." microsoft/USBroadbandUsagePercentages, GitHub repository; `https://github.com/microsoft/USBroadbandUsagePercentages`.

[6] U.S. Census Bureau, "Delineation files," August 4, 2023. Census.gov; `https://www.census.gov/geographies/reference-files/time-series/demo/metro-micro/delineation-files.html`.

[7] A. Riebler, S. H. Sørbye, D. Simpson, and H. Rue, "An intuitive Bayesian spatial model for disease mapping that accounts for scaling," *Statistical Methods in Medical Research*, vol. 25, pp. 1145–1165, August 2016. doi: 10.1177/0962280216660421.

[8] M. D. Hoffman and A. Gelman, "The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo," *Journal of Machine Learning Research*, vol. 15, no. 47, pp. 1593–1623, 2014. `https://jmlr.org/papers/v15/hoffman14a.html`.

[9] D. Phan, N. Pradhan, and M. Jankowiak, "Composable effects for flexible and accelerated probabilistic programming in NumPyro," *arXiv:1912.11554*, 2019. doi: 10.48550/arXiv.1912.11554; `https://arxiv.org/abs/1912.11554`.

[10] U.S. Census Bureau, "S2801: Types of computers and internet subscriptions," 2023. 2023 American Community Survey 5-Year Estimates, Subject Tables; `https://data.census.gov/table/ACSST5Y2023.S2801`.

[11] U.S. Bureau of Labor Statistics, "Quarterly census of employment and wages," 2024. https://www.bls.gov/cew/.

[12] O. Sidoti and C. McClain, "34% of U.S. adults have used ChatGPT, about double the share in 2023," June 25, 2025. Pew Research Center; https://www.pewresearch.org/short-reads/2025/06/25/34-of-us-adults-have-used-chatgpt-about-double-the-share-in-2023/.

[13] L. Saad, "AI: Major threat or just the next tech thing?," 2025. Gallup; https://news.gallup.com/poll/692435/major-threat-next-tech-thing.aspx.

[14] K. Ognyanova and V. Singh, "National AI opinion monitor: AI trust and knowledge in america," February 2025. Open Science Framework; doi: https://doi.org/10.31219/osf.io/ryv86_v1; https://naiom.net/public-reports/NAIOM%20Report%2002%20AI%20Trust%20Knowledge.pdf.

[15] A. Bick, A. Blandin, and D. J. Deming, "The rapid adoption of generative AI," Working Paper 32966, National Bureau of Economic Research, September 2024. doi: https://doi.org/10.3386/w32966; https://www.nber.org/papers/w32966.