Candidate Statement

Alexander Gamero-Garrido

My research is at the intersection of computer networks and public policy, with an emphasis on online privacy. My work yields methodologies and analyses that provide empirical evidence on contentious questions in technology policy, combining large-scale Internet measurements with social science methods.

Internet measurement infrastructure, which is used to collect data that informs high-stakes decisions by both governments and companies, was built by and for the Global North's core networks. My research asks what we get wrong when we treat that infrastructure as a neutral vantage point, and what becomes visible when we build methods that don't: which countries are exposed to state surveillance, which companies are tracking users and from where, and who has adequate connectivity at all. The thread running through this work is a core measurement problem: *widely-used vantage points and inference techniques systematically misrepresent the network, and the consequences of that misrepresentation are uneven across countries and populations*. My research spans three areas: *(i)* mapping and measuring the macroscopic structure of the Internet to reveal infrastructure-level surveillance capabilities and connectivity gaps; *(ii)* auditing commercial platforms and online tracking ecosystems to reveal what data is collected, where it flows, and whether regulatory frameworks are effective in practice; and *(iii)* characterizing Internet connectivity and performance at the institutions and communities that current measurement infrastructure systematically overlooks.

These research directions were initiated during a two-year postdoctoral fellowship at Northeastern University, which was primarily supported with independent funding from the Ford Foundation and the Northeastern Future Faculty Fellowship—both nationally competitive awards that enabled building a research agenda prior to joining UC Davis in 2023. My lab publishes primarily in the field's top venues for Internet measurement and online privacy: the ACM SIGCOMM Internet Measurement Conference (IMC) and the Privacy Enhancing Technologies Symposium (PETS), both of which are competitive, peer-reviewed venues that set the agenda for empirical research in their respective fields. Cross-disciplinary work appears in *PNAS Nexus*, a flagship open-access journal of the National Academies of Sciences, Engineering, and Medicine. A summary of publications, pre-prints, grants awarded, pending, and planned appears in Table 1; further subsections describe each research direction in detail.

During the review period, my lab's work has received the **IMC Best Paper Award** and the **Caspar Bowden PET Award Runner-Up** for the same paper; the **PAM Best Dataset Paper Award**; and the **CAMPOS Faculty Scholar** designation from UC Davis and the California State Legislature. **Total awarded funding** stands at $1,394,000, of which $632,000 has been directly awarded to me and my lab during the review period.

Research Area	Publications (in submission/resubmission)	Grants as PI/Co-PI & (Lab) Fellowships (Submitted/planned)
Computer Networks	†*PNAS Nexus* '23*, †*PAM '22* 🏆, †IMC '21 SIGCOMM '18 🏆 (IMC '26-c), (IMC '26-d) *lead computer scientist, interdisciplinary	3x †ISOC Fellowships [$44k total], Microsoft Dissertation Grant '19 [$25k] (2x NSF NeTS '27-30), (CAREER '27-32)
Consumer Privacy	†*IMC '25*, †*PETS '25*, †IMC '23 🏆🏆 (*IMC '26-a*), (*PETS '27*), (CCS '26)	†Ford Foundation Fellowship '22 [$50k], †NU Future Faculty Fellowship '21 [$100k] (NSF SaTC '27-30)
Internet Policy	*CCS '17*, *TPRC '15* (*IMC '26-b*), (*CCR '26-a*), (CCR '26-b)	†NSF NeTS '24-27 [$1.1M; $360k at UCD], †Comcast Innovation Fund [$85k], †CRA REU [$10k], (ISOC Research '26)

Table 1: Research areas with publication and funding records, with those in review period marked with †.
Italicized papers indicate those where I'm first author or senior lead (which, in computer science, is the last author.)
A 🏆 signals a field-wide distinction, such as the 2024 PETS Caspar Bowden Award Runner Up or the Best Paper Award at IMC 2023.

1 Research I: Computer Networks

My research in computer networks develops empirical methods to map and measure the Internet's physical and logical infrastructure, with a focus on making observable what is structurally hidden.

Published work. The foundation of this program is the Country-Level Transit Influence (CTI) metric, which I designed to quantify how much of a country's inbound international traffic is delivered by a single Autonomous System (AS)—the networks operated by ISPs, companies, and governments that collectively make up the Internet's routing fabric. The core methodological challenge is that BGP, the protocol through which networks exchange routing information, provides only partial and unevenly distributed visibility into the global routing system: monitors (data collectors) are concentrated in wealthy countries, backup paths are rarely visible, and large transit ASes span many countries in ways that make nationality assignment non-trivial. CTI addresses these limitations through a weighted metric that discounts indirect transit influence in proportion to topological distance, corrects for monitor clustering, and applies outlier filtering to remove monitor-location noise. CTI thus produces conservative *estimates of a country's exposure to traffic observation and selective tampering* despite incomplete data. Identifying which countries to study is itself a separate inference problem: a two-week active measurement campaign across 100 candidate countries, using 33 million traceroutes (network traces that reveal routers traversed between two endpoints), determined which nations rely primarily on transit rather than peering for international connectivity—the subset where CTI's inferences are most reliable. Applied to the resulting 75 transit-dominant nations, CTI reveals that in the median country, a single transit network sits on paths toward over 35% of the nation's IP addresses, creating a centralized point of exposure. The study (*PAM '22*, **Best Dataset Paper Award**) was validated through consultation with operators at 123 ASes across 19 countries. The study received mentions by the national (Financial Times) and international press.

That infrastructure-level dataset is foundational for two additional studies. The first paper (*IMC '21*) produced a global dataset of state-owned Internet operator ASes (revealing 989 ASes across 123 countries) through a multi-stage methodology crossing BGP data with corporate ownership records, regulatory filings, and investigative sources. The dataset revealed that in 18 countries, state-owned operators control over 90% of the estimated Internet access market, and that 19 countries operate Internet subsidiaries abroad, a phenomenon particularly concentrated in Africa. The second study (*PNAS Nexus '23*), a cross-disciplinary collaboration with political scientists at the University of Konstanz in which I served as lead computer scientist, combined both datasets in a novel empirical framework to ask a question neither field had previously been able to answer: does Internet topology differ systematically between autocracies and democracies? Prior research on autocratic Internet control focused on observable application-layer interference; this study reveals the structural conditions that make such interference possible without political backlash. Using statistical analyses and controlling for potential confounders (including development level and market structure), we show that in autocracies, state-owned providers hold significantly higher transit influence than private providers—a pattern absent in democracies—and that state-owned providers from autocratic countries operate preferentially in other autocracies, forming clusters of technological cooperation.

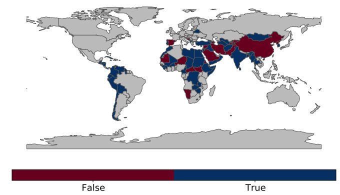

Transit-dominant nations (PAM).

Prior to starting my independent research program, I contributed to a highly visible study on inferring persistent interdomain congestion (*SIGCOMM '18*, **Best Paper Award**), which developed a measurement system deployed across 86 vantage points to study congestion on thousands of interdomain links across US broadband providers. This study provides empirical grounding for interconnection disputes at a scale previously unavailable.

Work in progress. Current work extends this infrastructure-mapping agenda in two directions. SAILOR (*IMC '26-c*, under review), developed through the Internet Society Pulse Research Fellowship (mentor track) and led by Riya Ponraj from the University of Oregon, introduces a hybrid learning-based framework for mapping traceroutes to submarine cable paths. This work addresses technical challenges on topological diversity, routing dynamics, and measurement noise. Separately, with my lab's incoming postdoc Sachin Kumar Singh, I am studying HTTP/3 adoption across 1.6 million domains including government sites in 30 countries (*IMC '26-d*, under review), revealing that while 35% of domains advertise HTTP/3, fewer than half publish the DNS resource record that enables clients to use it without a prior TCP handshake. The implication is that users pay a performance penalty on first connection that a single DNS record change would eliminate.

Nishant Acharya, an Internet Society Pulse Research Fellow in my lab, is developing inference-based methods to recover IXP topology from incomplete passive data and quantify how disruptions at these structural chokepoints propagate into measurement reliability. The core challenge is that passive BGP and traceroute data severely undersample IXP connectivity: in Sydney, CAIDA prefix probes recover only 3 bilateral peerings out of 135 EdgeIX members. The framework combines passive inference via metAScritic with targeted RIPE Atlas measurements to fill topology gaps, then simulates shutdown scenarios using BGP best-path selection to quantify impact on latency and path redundancy. Preliminary work across 12 countries with Cloudflare and ISOC found that active probing recovers 795 times as many links as traceroute alone, with 10% of links having no shorter alternative path under failure. An earlier version was presented at the IMC PhD Student Workshop.

Funding and pipeline. This research agenda is supported by three Internet Society Pulse Research Fellowships ($44k total). Two additional NSF NeTS proposals are in preparation: one investigating cloud infrastructure outage structure and cross-provider dependencies, and one examining speed tests as predictors of application-layer quality of experience across geospatial and wireless contexts. An NSF CAREER proposal advancing this agenda received three favorable panel reviews in its first submission—two rated Very Good, one Good—with the panel deliberating between Competitive and Low Competitive; a resubmission is planned for July 2026.

2 Research II: Consumer Privacy

My research in consumer privacy develops measurement frameworks to make commercial data collection empirically observable; it reveals what platforms collect, where data flows, and whether regulatory frameworks are effective in practice. The methodological thread is consistent: existing tools see only part of what is happening, and building a complete picture requires combining browser-level and network-layer observations from vantage points the field has systematically neglected.

Published work. Measuring effectiveness of data localization regulations (which dictate where web servers must be physically located), such as the EU's GDPR, presents a methodological challenge. BGP and DNS were not designed to enforce jurisdictional constraints, so verifying whether localization rules are observed in practice requires a method that does not exist off the shelf: no existing platform allows researchers to collect all the necessary information to determine where relevant servers are located. My data localization study (*PETS '25*) developed a novel framework by combining of browser measurements (to identify the full set of domains contacted during a user session) with data-plane measurements (to resolve the physical server IPs those domains use), deployed at scale across more than 1,000 networks in 19 EU countries. No existing platform supports both layers simultaneously at this geographic scope; the methodology required careful coordination across separate measurement systems to produce a representative sample of EU user traffic.

A further challenge was the geolocation of servers, for which we launched measurements from both the source (near the user) and the destination (near the server), along with rDNS records where available, to validate IP geolocation estimates. The methodology's precision was validated through a controlled experiment using CloudLab servers with known US locations and AWS servers in the EU, achieving a true positive rate of 85% and precision of 1.0—confirming that the approach correctly identifies non-compliant servers while avoiding false positives. The study finds that 2.2% of servers processing EU requests are located in non-adequate jurisdictions (the legal term for non-EU and not widely exempted countries), with significant variation across countries. Server location tracks infrastructure proximity rather than regulatory intent: trackers are significantly more likely to be located in non-adequate countries for users in Southern and Eastern Europe. The law is responding to a symptom while the cause, concentrated infrastructure ownership, remains outside its reach.

A fundamental obstacle in geographically expanding web tracking research is that nearly all prior work (including my own in PETS '25) collects measurements from centralized infrastructure in the Global North, producing findings that may not generalize to the regions where most Internet users live. The Gamma platform, led by Sachin Kumar Singh (then a PhD student at the University of Utah, now my lab's incoming postdoc), addresses this problem by running synchronized browser and IP-layer probes on the same device, achieving the cross-layer visibility that neither layer alone provides, from volunteer-operated hardware in under-measured regions. Temporal alignment of browser and network observations is non-trivial: requests must be attributed to specific server IPs without the source-device separation that controlled infrastructure provides. Deployed across 23 countries, Gamma reveals substantial geographic variation in tracker hosting: countries in sub-Saharan Africa exhibit high rates of non-local tracker embedding, while unexpected regions including Kenya and Malaysia emerge as significant hosting hubs. This study (*IMC '25*) is the first systematic cross-regional account of web tracking infrastructure from real user vantage points in the Global South.

My co-authored study on the Amazon Alexa ecosystem (*IMC '23*, **Best Paper Award**, **Caspar Bowden PET Award Runner-Up**) faced a different obstacle: auditing a closed commercial platform without manufacturer cooperation or direct access to its ad auction. The methodology intercepted network traffic and analyzed downstream advertising patterns—what ads users received, and at what bid prices—to infer what Amazon had learned about them from their voice interactions. This indirect inference approach, operating entirely from observable effects rather than platform internals, revealed that Amazon processes voice data to build user interest profiles and that third-party advertisers bid up to 30 times higher for certain profiles; most third-party skills fail to disclose their data practices. Results were presented at the Federal Trade Commission and cited as evidence in a class-action lawsuit against Amazon. The work was also covered by *The Verge*, *Chicago Tribune*, and *la Repubblica*, and was featured in a *Wired* interview on smart speaker privacy in 2026. The Alexa findings *prompted Amazon to update its privacy policy and disclose that Alexa interaction data is used for ad targeting.*

Work in progress. The Gamma and PETS lines continue with two papers that reframe VPNs as a measurement problem. Commercial VPNs are widely used as vantage points in studies of censorship, localization, and tracking; but if the VPN provider itself conditions what a measurement records, conclusions attributed to a country are partly artifacts of the instrument. Isolating this effect requires a design that holds country and capture time fixed while varying only the provider, a constraint that rules out most existing datasets. The first paper (*IMC '26-a*, under review) addresses the fundamental challenge of isolating provider-induced distortion from genuine geographic variation: if measurements through different VPNs in the same country produce different results, the difference must be attributable to the provider rather than the location. The study uses concurrent measurements across four major providers in 14 countries, holding country and capture time fixed, to decompose VPN-induced bias across three layers: vantage identity (whether the exit node's AS and IP are consistent with the claimed country), DNS behavior (whether the provider intercepts or redirects resolution regardless of client configuration), and replica selection (whether CDNs steer traffic differently based on the exit network's peering relationships). The second (*PETS '27*, under review) applies the same concurrent design to privacy-relevant surfaces—tracker identity, cookie counts, cookie lifetime—and finds that while which organizations track a user is stable across providers, how much they track and for how long is provider-conditioned on roughly one in four sites. Together these studies establish that a commercial VPN is not a neutral vantage point but an active factor shaping measurement outcomes, with direct implications for the validity of prior work that treated VPN-based measurements as transparent. An earlier version of this work received two weak accepts at ACM WWW, informing the current submission.

An additional paper (*CCS '26*, under review), co-advised with Zubair Shafiq and led by PhD student Muhammad Jazlan, addresses server-side tracking, where publisher-controlled intermediaries forward data to trackers without exposing direct client-to-tracker connections, defeating blocklist-based detection tools. SST-Guard exploits a semantic invariant: any compliant server-side Google Analytics implementation must collect and transmit the same information as its client-side counterpart regardless of endpoint obfuscation, enabling detection that is robust to infrastructure changes.

Funding and pipeline. This research agenda is supported by the Ford Foundation Postdoctoral Fellowship ($50k), the NU Future Faculty Fellowship ($100k), and an Internet Society Pulse Research Fellowship ($17k, awarded to Sachin Kumar Singh). An NSF Secure and Trustworthy Cyberspace (SaTC) proposal is in preparation, extending this agenda in two directions: a full-stack measurement study of US government websites at the state level—mapping tracker geography, third-party dependencies, and security posture across tens of thousands of domains—and a longitudinal evaluation of post-quantum TLS adoption, tracking where and how quantum-safe cryptography is actually being deployed across the web.

3 Research III: Internet Policy

My research in Internet policy applies network measurement to questions where empirical evidence has been absent or inaccessible: how US federal regulations shape the behavior of security researchers, how the Internet infrastructure serving anchor institutions compares to residential networks, and what socioeconomic factors drive broadband performance disparities at neighborhood granularity. The methodological challenge across all three directions is the same: the targets of study do not appear in standard measurement infrastructure and require purpose-built approaches to observe.

Prior work. My study on the legal risks facing security researchers (*CCS '17*) combined a direct canvassing experiment with a structured survey of vulnerability researchers, finding that legal challenges constitute a more significant deterrent to security audits than four alternative factors—operationalizing a previously anecdotal argument about regulatory chilling effects on empirical footing. An earlier study (*TPRC '15*) analyzed the FCC's Measuring Broadband America dataset to characterize residential broadband performance, establishing that performance degradation events result from causes on both access and content provider sides.

Work in progress. The centerpiece of the current period is the NSF NeTS grant ('24–27), which funds a program to map and characterize the Internet connectivity of US community anchor institutions (CAIs)—libraries, schools, healthcare facilities, and community centers—at a scale no prior work has attempted. The core methodological obstacle is that CAIs do not appear in any public, refereed database of network endpoints: they lease connectivity from ISPs and typically have no independently operated AS. My Reverse IP Geolocation (RG) framework, now substantially extended in a paper led by my lab's PhD student Nishant Acharya (*IMC '26-b*, under review), inverts commercial IP geolocation databases to infer IP ranges from physical addresses—combining reverse DNS, automated web crawling, NLP-based content classification, and active measurements. Applied to the full set of US public libraries across multiple states, RG identifies IP addresses serving institutions that are invisible to standard measurement infrastructure.

Two papers extend this agenda. First, *CCR '26-a* (planned resubmission in August), led by my PhD student Shivani Kalamadi, presents a framework for diagnosing broadband inequality at Census Block Group granularity—pairing 170 million crowdsourced Ookla speed tests with US Census demographic data across six metropolitan regions. A central methodological contribution is demonstrating empirically what is lost at coarser resolutions: without population-proportional bias correction and CBG-level granularity, analyses produce systematically different conclusions about which socioeconomic factors drive performance disparities. After controlling for population density, income and racial composition emerge as the operative drivers, with directions and magnitudes that differ sharply by region, ruling out a single national narrative. An earlier version received two positive reviews at TMA (accept and weak accept) and was presented at the IMC PhD Student Workshop (2025). Second, *CCR '26-b* (under review), led by collaborators at UC Santa Barbara, presents the first large-scale longitudinal analysis of broadband performance, stability, and reliability at 227 CAIs across seven states using nearly two million hourly speed test measurements from deployed Raspberry Pi devices (hosted by industry partner Exactly Labs).

Two additional directions are in progress: a study of path diversity between CAIs and cloud providers, led by collaborators at Northeastern University, and a congestion detection framework for externally monitoring CAI performance without client-side software or operator cooperation, led by Jiayi Liu (UC Santa Barbara).

Funding and pipeline. This research agenda is anchored by the NSF NeTS collaborative grant (*NSF NeTS '24–27*, $1.08M total, $360k at UCD, sole site PI), a collaboration with UCSB and Northeastern which I established during my postdoctoral fellowship, and supported by the Comcast Innovation Fund ($85k). A web-performance measurement tool developed with the Comcast grant is the foundation for the preliminary work included in a pending ISOC Foundation Research grant ($500k, $355k to UCD), which would extend this agenda by characterizing when and how speed test results predict application-layer quality of experience—across application types, access technologies, and deployment contexts including community anchor institutions—with particular focus on calibrating this mapping to historically under-measured regions where coverage is sparse and demographic data coarse.

4 Integration of Mentoring, Teaching and Research

A persistent gap in computer systems education is that courses emphasize engineering and implementation while Internet measurement demands something different: empirical reasoning that centers observation, inference, and statistical analysis, all grounded in real-world data. My teaching and mentorship program is designed around closing this gap, and it is deliberately integrated with the research agenda rather than parallel to it.

Mentorship model. My lab currently includes a postdoc (Sachin Kumar Singh), three solo-advised PhD students (Nishant Acharya, Shivani Kalamadi, Humaira Fasih Ahmed Hashmi), and one co-advised PhD student (Muhammad Jazlan, with Zubair Shafiq). Sachin's trajectory illustrates the lab's external reach: he sought me out as a collaborator while completing his PhD at the University of Utah, drawn by the lab's internet measurement expertise, and has since become first author on four submitted papers across IMC and PETS before joining as postdoc (start date in late June, 2026). Beyond Davis, I serve as a mentor in the Internet Society Pulse Fellowship program; Riya Ponraj (University of Oregon) and Karthik Nishanth (University of British Columbia) are external mentees, and Nishant Acharya received his own ISOC Pulse Fellowship, which *made my lab the only one in the 2025 cohort to hold both a mentor and a student fellowship simultaneously.* Within the lab, senior PhD students co-mentor undergraduate researchers, and also lead graduate code review and writing feedback sessions, building mentorship capacity across tiers. As a result of this broad mentoring model, my lab co-authored four IMC PhD student workshop papers in 2025 (over 20% of the program).

External service. I co-chaired the IMC 2024 Student Travel Grants, securing $19,000 in NSF funding and applying inclusive principles in allocation. More broadly in the field, I have been invited to serve on two NSF panels; I also serve on multiple program committees annually, especially for the top venues in my research fields (IMC, PETS).

Courses as research pipelines. I teach two project-based graduate courses at UC Davis: Computer Networks and Internet Measurement & Policy. The goal is for students to leave able to form hypotheses, collect measurements, and reason about uncertainty—skills that translate directly into research and industry practice. To further these goals, I structure projects around regular milestones with clear rubrics, have dedicated time to discuss project progress with instructional staff as part of the course schedule, require students to engage with measurement infrastructure directly (RIPE Atlas, M-Lab, Ookla), and ask them to interpret results in relation to both technical causes and external factors such as geography and policy. Both courses require original research projects and routinely produce work that matures beyond the classroom. For example, the geospatial disparities direction in my Internet Policy research began as a class project comparing broadband performance in Kansas, Arizona, and India, which later informed the bias-correction methodology in our planned resubmission (*CCR 26-a*).

5 Representative Papers and Pre-Prints

The papers listed below are a representative sample across my lab's three research areas. A full list of publications appears in my CV. Full papers are available in a shared Google Drive folder. **Directly* supervised student.

1. **Alexander Gamero-Garrido**, Esteban Carisimo, Shuai Hao, Bradley Huffaker, Alex C. Snoeren, Alberto Dainotti. Quantifying Nations' Exposure to Traffic Observation and Selective Tampering. *Passive and Active Measurement (PAM)*, 2022. **Best Dataset Paper Award.** [*PAM '22*]

2. Eda Keremoğlu, Nils B. Weidmann, **Alexander Gamero-Garrido**†, Esteban Carisimo, Alberto Dainotti, Alex C. Snoeren. Network Topology Facilitates Internet Traffic Control in Autocracies. *PNAS Nexus*, 2023. [*PNAS Nexus '23*]
†*Lead computer scientist, interdisciplinary team.*

3. Sachin Kumar Singh, Robert Ricci, **Alexander Gamero-Garrido**. Where in the World Are My Trackers? Mapping Web Tracking Flow Across Diverse Geographic Regions. *ACM Internet Measurement Conference (IMC)*, 2025. [*IMC '25*]

4. **Alexander Gamero-Garrido**, *Kicho Yu*, Sumukh Vasisht Shankar*, Sachin Kumar Singh*, Sindhya Balasubramanian*, Alexander Wilcox*,* David Choffnes. Empirically Measuring Data Localization in the EU. *Privacy Enhancing Technologies Symposium (PETS)*, 2025. [*PETS '25*]

5. *Sachin Kumar Singh**, Robert Ricci, **Alexander Gamero-Garrido**. Not All Roads Lead to Rome: How VPN Selection Alters What We Measure and Infer about Web Infrastructure. *In submission, ACM Internet Measurement Conference (IMC)*, 2026. [*IMC '26-a*]

6. *Nishant Acharya**, Anyu Yang, *Humaira Fasih Ahmed Hashmi**, Kevin Vermeulen, *Shivani Kalamadi**, Jiayi Liu, *Ashutosh Kshirsagar**, Elizabeth Belding, David Choffnes, **Alexander Gamero-Garrido**. Where's Waldo Library? Using Reverse IP Geolocation to Identify Library IPs. *In submission, ACM Internet Measurement Conference (IMC)*, 2026. [*IMC '26-b*]