

Auditing Alignment Controllability in LLMs via Political Axes

Bartol Bućan^{1*}, Nikola Soćec¹, Sarah Isuff², Morena Granić¹, Luka Hobor¹, Agneza Krajna¹,
Mihael Kovac¹, Mario Brcic^{1†}

¹ University of Zagreb Faculty of Electrical Engineering and Computing, Zagreb, Croatia,

² It From Bit d.o.o., Zagreb, Croatia

Abstract

Political audits of large language models (LLMs) usually reduce each to one point on a political compass. But that resting point barely matters in deployment: a model must land somewhere, and what counts is how far, and in which directions, its answers can be steered. That steering runs through the system prompt: the personalization layer a platform sets, or one induced from a user’s history, not necessarily written by hand. We run a dispersion-first stress test of prompt-based controllability across 12 ideological personas plus an unsteered baseline, 70 Political Compass items, ten replicates, and seven leading LLMs: GPT-5, Claude, Grok, Gemini, DeepSeek, Kimi, and Qwen (63,700 responses). Contextual framing explains roughly 88%–93% of variance on the economic and society axes, model identity under 3%: responses are highly instruction-adjustable. Models do not shift alike: some move more, and some saturate under extreme framings. Conflicting directional-steering results in prior audits resolve once baselines are recognized as non-centered: displacement and proximity diverge, so the effect is geometric, not differential compliance. Under authoritarian prompts, models produce similar shifts on the same questions. Political-coordinate audits therefore need steerability audits reporting dispersion, symmetry, saturation, and refusal floors. We release prompts, benchmark data, and code.

Introduction

The standard approach to political evaluation places a model on a political compass and reports a coordinate, but a coordinate tells us where a model sits in the absence of pressure, not what happens when a user pushes. Two models with identical baselines may behave very differently in practice: one holding its ground under its inherent ideological framing, the other drifting with the system prompt.

A substantial body of work has established the baselines. Static audits consistently find that most LLMs cluster in the left-libertarian quadrant (Rozado 2023, 2024; Hartmann, Schwenzow, and Witte 2023; Peng et al. 2026; Sakhawat et al. 2026). But methodological critiques have shown that

these position estimates are surprisingly fragile: prompt wording, answer format, and paraphrase variation can materially alter inferred coordinates (Röttger et al. 2024a; Ceron et al. 2024). If the measurement itself is unstable, then the object being measured may not be a fixed property of the model at all.

Persona-based studies report that models shift ideologically under conditioning, with larger models showing broader coverage and asymmetric responsiveness (Bernardelle et al. 2025b,a). Aldahoul et al. observe that apparently moderate aggregate positions can mask offsetting extremes on specific topics, framing this as “ideological inconsistency” (Aldahoul et al. 2025). These findings suggest that political bias in LLMs is not a point but a distribution, and that the distribution’s shape matters as much as its center.

We study in-context political-axis steering not because Political Compass is a complete theory of values, but because it provides a controlled stress test for a more general prerequisite of pluralistic alignment: whether authorized instructions can move value-sensitive behavior predictably, symmetrically, and visibly. Controllability and value alignment are themselves subject to known theoretical limits (Brcic and Yampolskiy 2023); our contribution is to characterize them behaviorally for deployed LLMs. We define *ideological dispersion* as the average displacement of a model’s political outputs from its baseline under systematic thematic framing, and we treat it as the primary evaluation metric rather than a secondary observation. That connects to the emerging literature on steerable pluralism (Sorensen et al. 2024; Kirk et al. 2024) and to persuasion-risk evidence showing that even brief interaction with opinionated AI can shift user attitudes (Jakesch et al. 2023; Hackenburg and Margetts 2024; Hackenburg et al. 2025; Potter et al. 2024).

Contributions. (i) *Conceptual*. Political coordinates are single points on a surface of behavior that instructions move, not complete descriptions of deployed model behavior. (ii) *Empirical*. Within our forced-choice benchmark, system-prompt framing accounts for roughly 88% of economic-axis variance and 93% of society-axis variance while differences between models account for under 3% on both. (iii) *Diagnostic*. Controllability is non-uniform across dispersion, directional symmetry, saturation, refusal floors, and metric sensitivity; we name *metric non-equivalence un-*

*Bartol Bućan, Nikola Soćec, and Morena Granić were students at the University of Zagreb Faculty of Electrical Engineering and Computing at the time of this work.

†Corresponding author: mario.brcic@fer.hr
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

der non-centered baselines as a concrete audit artifact (displacement vs. proximity can rank directions differently). (iv) *Comparative*. Cross-model convergence in per-question shift patterns under authoritarian framing (aggregate Spearman $r = 0.79$, permutation $p < 0.001$), with the underlying system-layer driver explicitly left unresolved by black-box design.

Instruction-following alone predicts only *that* a role-conditioned model will move, not how far, whether the movement is symmetric, whether it saturates or reverses under pressure, or whether independently trained models converge on the same item-level pattern; those are the quantities we measure. Our design targets role-based, agentic deployments (Tseng et al. 2024; Hu et al. 2024) instead of open dialogue, and does not assume the user writes the system prompt by hand: a profile can be induced from prior sessions and refined by automatic prompt optimization (Yuksekgonul et al. 2025; Khattab et al. 2024; Zhang et al. 2024). We characterize that starting point and the distance a system can travel from it, deferring the boundary against dialogue-emergent controllability to § Limitations.

Related Work

Static Position Audits and Their Limits

The dominant paradigm in LLM political evaluation is questionnaire-based auditing. Rozado’s foundational work administered 15 political orientation tests to ChatGPT, finding left-leaning preferences in 14 of 15 (Rozado 2023), and later extended this to 24 conversational LLMs, most of which are diagnosed as left-of-center on both the economic and the social axis (Rozado 2024). Hartmann et al. used 630 political statements from voting advice applications and found a consistent pro-environmental, left-libertarian orientation (Hartmann, Schwenzow, and Witte 2023). The largest cross-model comparison to date covers 43 models from 19 families on position alone (Peng et al. 2026), while Sakhawat et al. audit 26 models across three psychometric inventories and find most of them clustered in a single ideological region, with 96.3% of placements in the Libertarian-Left quadrant (Sakhawat et al. 2026). Their variance decomposition holds the persona fixed and varies instrument wording, where model identity dominates ($\eta^2 > 0.90$). Ours holds wording fixed and varies the persona, which is why the two decompositions assign the variance in opposite directions.

These studies establish important baselines, but they treat political orientation as a point estimate that is surprisingly fragile: format variation and paraphrase can materially alter inferred coordinates (Röttger et al. 2024a), and large models show inconsistency across topics when responses are sampled repeatedly (Ceron et al. 2024). If the position moves due to measurement variation, it is the movement, not the position, that characterizes the model.

Steerability, Dispersion, and Distributional Opinion Work

Santurkar et al. show that RLHF-tuned models disproportionately align with liberal, high-income, well-educated de-

mographics and, critically, that demographic steering is surprisingly ineffective at correcting this misalignment (Santurkar et al. 2023). This suggests that steerability is not uniform across the ideological landscape.

Persona-based studies directly observe ideological shifts. Bernardelle et al. test seven models with synthetic persona biographies and find that larger models show broader ideological coverage, that susceptibility to ideological cues grows with scale, and that responsiveness is asymmetric, stronger toward right-authoritarian than left-libertarian positions (Bernardelle et al. 2025b,a). Aldahoul et al. frame related findings as “ideological inconsistency,” demonstrating that LLMs have lower variance than voters but higher variance than legislators (Aldahoul et al. 2025). Batzner et al. have six commercial LLMs role-play the personas of German parliamentary group leaders against 418 voting-advice-application statements, and argue that the resulting shift toward the prompted party’s position reflects persona-based steerability rather than the increasingly popular, yet contested, concept of sycophancy (Batzner et al. 2025).

Formal steerability benchmarks complement this empirical work. The NAACL 2025 steerability benchmark defines steering as a distribution shift relative to the baseline and proposes steerability indices (Miehling et al. 2025). Li et al. improve steerability using collaborative filtering to embed opinions into continuous vector spaces (Li et al. 2024). Chang et al. reveal that recent LLMs are less steerable than they appear, attributing this to “side effects”: correlations between requested and unrequested behavioral changes (Chang et al. 2025); our compass-only audit cannot detect non-compass-axis side effects and we flag this as a scope constraint (§ Limitations).

Closest to our study, Bernardelle et al. examine persona-induced shifts on Political-Compass-style instruments using synthetic persona biographies on open models (Bernardelle et al. 2025b,a), and Batzner et al. measure persona-role-play-induced shifts toward prompted-party positions on six commercial LLMs against a German, party-referenced instrument (Wahl-o-Mat) (Batzner et al. 2025). Our contribution differs in four respects. First, we treat political axes as a graded controllability stress test (twelve framings at three intensity levels, plus a baseline), using directly ideologically-loaded statements whose correct handling does not depend on country- or party-specific platform knowledge. This avoids the single-intensity role-play against party-position ground truth used by Batzner et al., where reproducing a party’s platform is itself a confound they report for their own instrument (particularly for centrist parties). Second, we evaluate seven leading commercial frontier endpoints on a four-axis, English-language political-compass instrument, not a single-country party system or open-source model families. Third, we decompose within-benchmark variance across framing and model factors and report metric non-equivalence between displacement and proximity, going beyond a single relative-shift measure. Fourth, we report cross-model per-question shift convergence as an audit diagnostic, leaving the underlying system-layer driver unresolved by black-box design. Two further studies frame our closest neighbours. Smith-Vaniz et al. probe political and

demographic leanings through Moral Foundations Theory, comparing inherent, explicitly-prompted, and demographic-persona conditions against human MFT survey data (Smith-Vaniz et al. 2025). Like us, they use persona role-play to elicit political leaning. Unlike us, they benchmark ideological *accuracy* against that human data; we decompose framing-versus-model variance and characterize the controllability *profile* (dispersion, symmetry, saturation once stronger framing stops adding movement, and the refusal rate framing cannot push below) rather than positional correctness.

Mechanisms, Governance, and Deployment Context

Understanding *why* models differ in dispersion requires engaging with both mechanistic and policy-level work. On the mechanistic side, high dispersion is linked to sycophancy, the tendency of models to mirror stated user views regardless of content (Sharma et al. 2024; Perez et al. 2023). Low dispersion, conversely, may reflect over-refusal or rigid safety regimes that avoid engagement with politically sensitive topics (Cui et al. 2025; Röttger et al. 2024b). Interpretability work suggests a deeper structural basis. Kim et al. demonstrate that political ideology is encoded as linear representations in activation space, functioning like a tunable dial (Kim, Evans, and Schein 2025). Cintas et al. localize it further: persona representations for conservatism and liberalism occupy more distinct regions of activation space than competing ethical-framework personas, and concentrate in the final third of decoder layers (Cintas et al. 2025). Kabir introduces “ideological depth,” showing that some models possess far richer political feature representations than others, and that refusals often stem from capability gaps, not safety rules (Kabir 2025). For Chinese-origin models, geopolitical censorship patterns, documented on questions of Taiwanese sovereignty, may also contribute to low-dispersion profiles through hard content filtering, not alignment training (Ko 2026).

On the governance side, industry evaluations now assess political behavior directly. OpenAI’s five-axis framework measures political refusals, asymmetric coverage, and emotional escalation (OpenAI 2025). Anthropic’s evenhandedness evaluation compares paired-prompt responses across multiple frontier models, including several in our study (Anthropic 2025). These evaluations provide essential context: the dispersion patterns we observe are not just technical properties of models but consequences of deliberate alignment choices by their developers.

Methodology

Experimental Design

We designed our study to isolate the effect of ideological prompt framing on model outputs, holding all other variables constant. We evaluated seven frontier models across 13 conditions (12 systematically varied ideological framings plus one unsteered baseline) with 70 questions per condition and a balanced core of ten replicates per model-condition cell. This yields $7 \times 13 \times 70 \times 10 = 63,700$ question-level

responses. We accordingly report directional tier groupings rather than strict ordinal rankings, and provide bootstrap CIs throughout to make remaining uncertainty explicit.

Question Instrument

We use a 70-item multi-axis political questionnaire in the Political Compass/8values tradition of public political-orientation quizzes, scoring responses across four axes: economic (left–market), diplomatic (world–nation), government (liberty–authority), and society (progress–tradition). We acknowledge the well-documented sensitivity of Political Compass–style instruments to format and phrasing (Röttger et al. 2024a); our dispersion metrics are relative within-instrument comparisons and do not depend on the absolute accuracy of any single position coordinate. Each item is answered using one of five labels (Strongly Disagree through Strongly Agree), mapped to ordinal scores in $\{-2, -1, 0, 1, 2\}$ and aggregated with per-item axis effects into normalized percentage scales. The five-label scale and the $\{-2, \dots, 2\}$ mapping are inherited from the instrument, not introduced here, and match the scoring convention used by prior LLM political audits (Rozado 2024; Hartmann, Schwenzow, and Witte 2023). Because dispersion, displacement, and proximity are all distances in the resulting coordinate space, they are invariant to affine rescaling of the ordinal map: unit spacing affects absolute coordinates, which we do not treat as load-bearing, not relative movement. Our primary analysis focuses on the economic and society axes because (econ, society) is the display plane used by the static audits we compare against (Rozado 2023; Hartmann, Schwenzow, and Witte 2023), which keeps our baselines commensurable with theirs. It is also the conservative choice: the government axis, which our 0A/0L framings target most directly, shows an even larger context effect (§ Results), so reporting society as primary understates rather than inflates the headline result. Government and diplomatic axes are used for robustness checks and for the signed-axis analysis.

Context Taxonomy

The 12 steering conditions vary along two dimensions: an ideological axis and an intensity level (Table 1). Context codes follow the convention [econ][social]-[intensity], where the economic position is L (left), R (right), or 0 (neutral); the social position is A (authoritarian), L (libertarian), or 0 (neutral); and intensity runs from _1 (moderate) through _3 (radical/hardline). Thus 0A_3 denotes a socially authoritarian persona at maximum intensity with no economic tilt, while L0_3 denotes a radical economic-left persona with no social axis bias.

Steerability is operationalised through *system-prompt persona injection*. Each condition replaces the default system prompt with a paragraph-length ideological persona description (70–250 words), followed by the standard response-format instruction. The persona is written in second person (“You believe...”; “You hold that...”) to prime the model to respond *as* that persona rather than *about* it. For example, the 0A_3 condition (*The Hardline Social Authoritarian*)

Code	Axis	Intensity
LO_1..3	Economic left	Moderate → radical
RO_1..3	Economic right	Moderate → hardline
OL_1..3	Social libertarian	Moderate → radical
OA_1..3	Social authoritarian	Moderate → hardline
None	Baseline	No framing

Table 1: Context taxonomy for ideological prompt framing.

opens: “Essence: Order sanctified, control as salvation. Sees human nature as dangerous, sinful, and weak...” The LO_3 condition (*The Revolutionary Economic Leftist*) frames all wealth as structural theft and collective ownership as the only moral basis for production. Intensity levels modulate the rhetorical register: _1 framings use hedged, moderate language; _3 framings use maximalist, ideologically committed language. Across ideological *directions* the personas share a common template, length band, and second-person register, differing only in ideological content, so that direction and intensity, not surface style, drive the contrasts; a length-matched neutral control (§ Limitations) confirms that prompt length is not the driver. All 12 persona texts and the baseline system prompt are released in the reproducibility package.

Models and Inference Settings

All seven models were accessed via commercial API routing over two days (2026-05-19 to 2026-05-21, UTC). Inference settings were held constant: temperature = 0.7, forced single-label output. Temperature 0.7 is a deliberately stochastic conversational setting rather than a near-deterministic evaluation setting. Testing at $T = 0.7$ is therefore the harder test for robustness: a steerability effect that survives sampling noise is more credible than one that only appears at $T \approx 0$. A temperature ablation at $T = 0.1$ is reported in § Limitations. Model API identifiers are: openrouter/openai/gpt-5, openrouter/anthropic/claude-sonnet-4.5, openrouter/google/gemini-2.5-flash-lite, openrouter/x-ai/grok-4.3, openrouter/deepseek/deepseek-chat-v3.1, openrouter/moonshotai/kimi-k2-0905, and openrouter/qwen/qwen3.6-max-preview.

Metrics

We define three families of metrics. Let $\mu_{m,c,r} = (\text{econ}, \text{scty})$ be the 2D centroid for model m , context c , and run r .

Dispersion from baseline. The primary metric is the average Euclidean distance between steered and unsteered cell centroids. Let $\bar{\mu}_{m,c} = \frac{1}{|R|} \sum_{r \in R} \mu_{m,c,r}$ be the cell centroid for model m in context c averaged across replicates. Then

$$D_m = \frac{1}{|C^*|} \sum_{c \in C^*} \|\bar{\mu}_{m,c} - \bar{\mu}_{m,\text{None}}\|_2,$$

where C^* excludes the baseline condition; a higher D_m means the model moves more under ideological framing.

Asymmetry decomposition. To avoid misleading directional claims, we use two complementary measures for each steered condition: *displacement* (how far the model moved from its baseline, $d = \|\mu_{m,c,r} - \mu_{m,\text{None},r}\|_2$) and *proximity* (how close it ended up to the relevant display-plane extreme; lower = closer to target). For the Figure 2 projection, the targets follow the scoring convention: LO_3 → econ = 100, RO_3 → econ = 0, OL_3 → scty = 100, OA_3 → scty = 0. We separately check signed Authority/Liberty movement on the native government axis below.

Multiple-comparison awareness. The paper separates inferential checks (one tier Mann–Whitney contrast, four cross-model permutation tests, and three DeepSeek-only ablation checks) from descriptive diagnostics such as saturation counts and the per-intensity convergence-ladder $\bar{r}$ values. Large effects ($\eta^2 \approx 0.88\text{--}0.93$ on the primary axes, $\bar{r} \approx 0.8$) survive any reasonable multiplicity correction; borderline diagnostics are flagged as exploratory rather than confirmatory.

Missing-data policy. Non-Likert, refusal-like, empty, and API-failed responses are not treated as Neutral. We preserve the raw output, log the failure type (refusal, unparseable, empty, api_error), and exclude missing responses from both numerator and axis-denominator in score aggregation. A single pre-specified repair retry using the *identical* elicitation prompt as the original call is applied as a supplementary analysis; persistent failures remain missing. Repair counts and the resulting bound on headline aggregates are reported in § Limitations.

Results

What You Prompt Matters More Than Which Model You Use

The most striking finding is quantitative: context framing overwhelmingly dominates inter-model baseline differences in explaining observed ideological variance. Additive variance decomposition on valid responses yields:

- Economic axis: $\eta_{\text{ctx}} = 0.8816$, $\eta_{\text{model}} = 0.0250$.
- Society axis: $\eta_{\text{ctx}} = 0.9313$, $\eta_{\text{model}} = 0.0072$.

Per-axis, the government axis is the most context-driven ($\eta_{\text{ctx}} = 0.9558$), followed by society (0.9313), diplomatic (0.9182), and economic (0.8816); model main effects on every axis are under 3%.

High η_{ctx} is partly by construction: our contexts were designed to span the ideological space. The low η_{model} is not. Our design could have failed to find it, because a model with rigid alignment guardrails or a refusal floor would have produced flat dispersion and a substantial η_{model} . Refusals stay at 1.19% even when the prompt explicitly permits them (§ Limitations), and neutral responses all but vanish under extreme framing (§ Saturation). The inter-model differences that surface in static audits are small relative to within-model movement under system-prompt framing. Adding a context × model interaction term to the same decomposition leaves the main effects essentially unchanged but absorbs $\eta_{\text{int}} = 0.0837$ on the economic axis and $\eta_{\text{int}} = 0.0508$ on the society axis out of the residual. The interaction is therefore 3–7× larger than the model main effect: model identity

matters less as a static baseline offset and more through how each model responds to each framing. The tier separation, per-model saturation pattern, and cross-model coordination findings reported in the rest of this section all sit inside this interaction term rather than inside η_{model} .

Models Separate into Dispersion Tiers

Despite the dominance of context, models do differ in how much they move. Mean dispersion scores, from highest to lowest, are: Kimi K2 (33.32), Qwen3.6 Max Preview (33.28), GPT-5 (31.41), Claude Sonnet 4.5 (29.17), Grok-4.3 (28.75), Gemini 2.5 Flash Lite (26.49), and DeepSeek-Chat v3.1 (24.04). A one-sided Mann–Whitney on the top-three/bottom-four split gives $U = 12$, $p = 0.029$ (post hoc; reported descriptively, not as confirmatory inference); the high-tier mean (32.67) exceeds the low-tier (27.11) by $1.21\times$.

Confidence intervals for dispersion are computed by context-level bootstrap over the 12 steered cell-centroid distances within each model, with $B = 5,000$ resamples and percentile intervals; 95% intervals are broad and adjacent models' intervals overlap substantially. Across $n=10$ replicates per cell the data support grouping into tiers, not a strict ordinal ranking. We identify two tiers: a higher-dispersion group (Kimi K2, Qwen3.6 Max Preview, GPT-5) that shows approximately 20% more movement from baseline, and a lower-dispersion group (Claude Sonnet 4.5, Grok-4.3, Gemini 2.5 Flash Lite, DeepSeek-Chat v3.1) that maintains more stable baselines across conditions.

The tier composition is largely stable across all four axes. Kimi K2 and Qwen3.6 Max Preview are the top-two highest-dispersion models on every axis (economic, society, government, diplomatic). Gemini 2.5 Flash Lite and DeepSeek-Chat v3.1 are the bottom-two on the society, government and diplomatic axes; the economic axis is the exception, where DeepSeek-Chat v3.1 is still lowest but Gemini 2.5 Flash Lite rises to third. Direction agrees across axes but magnitude does not (Pearson $r = 0.54$, $p = 0.21$, $n = 7$, between per-model econ- and scity-axis dispersion). Gemini 2.5 Flash Lite illustrates the magnitude uncoupling: 22.1 econ-axis dispersion vs. 12.3 scity-axis (a $1.8\times$ ratio). The tiering is not a 2D-aggregation artifact (Fig. 1); it reflects an axis-stable behavioral disposition whose strength varies by axis.

The Asymmetry Paradox

A naïve reading of our data produces a contradiction. Models appear to be more steerable toward right-economic and authoritarian positions: mean displacement is 51.97 for RO_3 versus 37.76 for LO_3, and 49.30 for OA_3 versus 25.50 for OL_3. Proximity gives a more axis-dependent story. On the economic axis, it reverses the displacement ranking: LO_3 ends closer to its display-plane target than RO_3 (7.85 vs. 11.48). On the society display axis, OA_3 also ends closer than OL_3 (17.09 vs. 27.27). Directional asymmetry therefore cannot be summarized by one scalar; it depends on whether the audit asks how far the model moved or how close it came to a specified target.

The 2D compass projection (Fig. 2) provides spatial context for these results. The data show that instruction-induced ideological movement is not axis-independent. Models respond to directional framings through culturally bundled ideological profiles, producing dense Left-Progressive and Traditional-Right-adjacent regions while leaving the off-diagonal Left-Traditional and Right-Progressive quadrants sparsely populated. Two complementary compound-quadrant pilots, in which contexts request both axes simultaneously, confirm that this direction-dependence is not an artifact of the single-axis prompt families: the same bundled diagonal is reached preferentially even under explicit off-diagonal prompts (§ Limitations; supplementary appendix).

The resolution is geometric. All seven frontier endpoints start in the Progress half of the social axis; six of seven start in the (Left, Progress) quadrant of the political compass; only Grok-4.3 deviates, and only on the economic axis (econ = 45.8, scity = 59.1). No frontier endpoint starts in either Tradition quadrant. Rightward and authoritarian prompts must therefore traverse a greater ideological distance, producing larger displacement values. Because the models start in the Left-Progress region, economic-left prompts need less movement to reach their display-plane target, while authoritarian prompts produce a larger traverse from a Progress baseline. Displacement measures effort; proximity measures attainment. Reporting only one produces a misleading picture: studies that report only “models are more steerable toward the right” may be describing baseline geometry rather than differential compliance.

A signed-axis check confirms the geometric reading. Under authoritarian framing (OA_3), mean shift on the government axis is -49.9 percentage points (all seven models cross the neutral midpoint into Authority territory; baseline mean $\approx 56 \rightarrow$ steered mean ≈ 6). Under libertarian framing (OL_3), the corresponding shift is $+29.3$ points. The asymmetry is real on the signed axis, but it is asymmetry-of-magnitude, not asymmetry-of-compliance: all framings produce shifts in the expected direction. No model exhibits an authoritarian-direction flat-response floor of the kind prior work has sometimes implied.

Saturation at Ideological Extremes

We observe a non-monotonic pattern in economic-left framing: past a point, stronger framing stops producing stronger movement. It complicates simple accounts of model behavior. Headline 3-of-7 on 2D L2 displacement: Gemini 2.5 Flash Lite, GPT-5, and Grok-4.3 produce *less* 2D displacement under the most extreme left-economic framing (LO_3) than under the moderate version (LO_2). On the directly targeted economic axis, 5-of-7: Gemini, Kimi K2, GPT-5, Qwen3.6 Max Preview, and Grok-4.3 score lower on LO_3 than LO_2. The two metrics disagree on Kimi K2 and Qwen3.6 Max Preview: their 2D displacement keeps increasing because the society-axis component continues to move under LO_3 even as the economic-axis component retreats. Radical prompting overshoots in either reading: it triggers a saturation pattern under which additional ideological pressure yields diminishing or partially reversed returns.

This finding is inconsistent with a pure sycophancy expla-




Figure 1: Steerability overview. Values are percentage-point distances on the normalized 0–100 Political Compass scales: the left panel reports each model’s mean 2D Euclidean displacement from its baseline across the 12 steered contexts, while the right panel reports mean absolute displacement on the economic and society axes separately.

nation. A model that simply mirrors user intent would comply uniformly with intensity: more extreme prompts should always produce more extreme outputs. The descriptive endpoint non-monotonicity is compatible with, but not diagnostic of, alignment guardrails activating at extreme steering intensities, or with diminishing marginal returns as models are pushed further into their own trained region. A soft-refusal alternative (that models hedge with neutral answers when pushed) is also inconsistent with our data: under LO_3 framing models pick the neutral answer in 1.92% of question-responses, compared to 37.93% at baseline; under authoritarian and libertarian extremes the rates fall below 1%. Models commit to non-neutral answers under framing; they do not hedge.

Cross-Model Coordinated Response Patterns

The dispersion patterns we report so far are model-by-model. Are the seven models, trained independently by seven labs, responding to identical framings in idiosyncratic ways, or in similar ones? For each (model, extreme-framing) cell we compute a per-question shift vector $\Delta_q = \bar{s}_{m,c,q} - \bar{s}_{m,\text{None},q}$ over the 70 PC items, then take the Spearman correlation of these vectors between every model pair. The mean pairwise rank correlation is striking: $\bar{r} = 0.79$ under authoritarian framing (OA_3), 0.81 under right-economic, 0.73 under left-economic, 0.66 under social-libertarian. A permutation test that shuffles each model’s shift vector independently and recomputes the mean pair correlation places

all four observed values far outside the null distribution ($p < 0.001$ in 2,000 shuffles; null 95th percentile ≈ 0.05).

In plain terms: seven frontier models trained by seven labs not only move under framing, they move in correlated, question-specific ways at the level of observed outputs. Under OA_3 , approximately three-quarters of the 70 PC items receive unanimous-sign agreement across all seven models under a strict nonzero-sign convention.

The convergence *scales with prompt intensity*: mean pairwise $\bar{r}$ on the per-question signed-shift vectors climbs monotonically with framing strength across all four ideological families, from the 0.34–0.46 range at moderate intensity to 0.66–0.81 at maximal intensity, strictly increasing at every step on all four (e.g. $0.46 \rightarrow 0.77 \rightarrow 0.79$ for OA). At moderate framing each model behaves idiosyncratically; under maximalist framing they converge toward a common item-level pattern, as a trained-range account would predict.

We treat this convergence as *output-level regularity in shift patterns*, not evidence of shared internal representations or post-training mechanisms: a black-box design cannot distinguish shared pretraining-corpus priors, instruction-following competence, RLHF preference conventions, or distillation effects. The shuffle null treats items as exchangeable; item dependence is a known limitation, and a within-axis block-permutation null is left to future work.




Figure 2: Political compass placements across all 13 conditions. Each cluster represents one model’s responses with context labels and 95% replicate ellipses computed from replicate-level (econ, scty) centroids. Context families separate in expected directions, but the occupied region is not a uniform grid: directional framings induce coupled movement across axes, with model-specific footprint sizes.

Discussion

Behavioral Predictions of Mechanistic Accounts

Interpretability work provides a structural prediction against which our behavioral results can be checked. Kim et al. (2025) show that political ideology is encoded as approximately linear directions in activation space, functioning like a tunable dial; Kabir characterises variation in “ideological depth” across models and finds that low-depth models often produce refusals from capability gaps, not safety rules (Kabir 2025). A linear-representation account makes two behavioral predictions: (i) when a framing pushes activation past a learned range, further pressure should yield diminishing or reversed effects (a saturation ceiling), and (ii) models with similar training pipelines should activate similar question-level features under identical framings, producing correlated per-question shift vectors. The saturation pattern and cross-model coordination reported above are consistent with both. They are equally the pattern a pure sycophancy account does *not* predict: a model that merely mirrors user intent should comply uniformly with framing in-

tensity rather than retreat at LO_3. None of this is diagnostic. Formal saturation tests are sensitive to the unit of analysis, so we treat the counts as descriptive endpoint non-monotonicity, not proven per-model reversal, and prompt-pathology and questionnaire-saturation explanations remain open under a black-box design. We do not claim our data identifies the mechanism; we claim it is the behavioral signature mechanistic theories of LLM ideology would predict, and offer it as a target for future interpretability work to explain or falsify.

From Position to Profile

Our results argue for a shift in how we evaluate the political behavior of LLMs. Context dominates inter-model baseline differences. A model’s static political coordinate is not meaningless, but on the primary axes it captures under 3% of the variance that matters in interactive deployment. A further 5–8% sits in the context $\times$ model interaction: model identity matters chiefly through how each model responds to framing, not through its default coordinate. The dispersion profile (how far a model moves, in which directions, and with what

reliability) is far more informative. A baseline offset and a steerability-profile limit are also not equally consequential. A baseline can be relocated by whoever sets the instruction layer; a reachability gap, saturation ceiling, or refusal floor (§ Limitations) is a property of the model that prompting cannot remove; correcting it requires finetuning, activation-level intervention, or retraining. A bias in *what the profile can reach* is therefore the more serious and less tractable failure than a bias in *where the profile starts*.

That does not make static audits obsolete. Baseline position determines the starting point from which displacement and proximity are measured, and the asymmetry paradox we document arises precisely because baselines are not centered. But it does mean that evaluations that report only position are answering a question that is less relevant than it appears for the settings where LLMs are actually deployed: conversations, tutoring, writing assistance, and information retrieval.

The Bounded Pluralism Dilemma

High dispersion is not intrinsically good or bad: its valence depends on what is driving it and where it is deployed.

Under a bounded-pluralism lens (Sorensen et al. 2024; Kirk et al. 2024), high dispersion can be desirable: it may indicate that a model can genuinely engage with diverse value systems, supporting users across the political spectrum instead of imposing a single ideological default. That is the promise of steerable pluralism: AI that adapts to its user’s values rather than its developer’s. Recent work operationalizes this directly: Adams et al. steer models across pluralistic value profiles via few-shot comparative regression (Adams et al. 2025). Our contribution is complementary and diagnostic, not prescriptive: we do not propose a steering method but measure how wide the reachable value range already is under ordinary system-prompt control, and where it saturates.

But high dispersion can also reflect sycophancy (Sharma et al. 2024; Perez et al. 2023): a model that shifts toward whatever position it detects in the prompt, not because it possesses genuine ideological range but because it has learned that agreement is rewarded. In this case, dispersion measures not pluralism but compliance.

The non-monotonic saturation pattern is the main evidence against a pure-sycophancy reading of high dispersion: it is compatible with, though it does not establish, an internal constraint such as an alignment guardrail or a trained ceiling on extreme-direction outputs. We develop that argument, and its limits, in § Behavioral Predictions above.

Low dispersion, conversely, can indicate principled consistency or rigid refusal to engage. Over-refusal benchmarks document the cost of safety alignment in helpfulness loss (Cui et al. 2025; Röttger et al. 2024b). For DeepSeek-Chat v3.1, which shows the lowest dispersion in our study, an additional factor may be relevant: geopolitical content filtering has been documented in Chinese-origin models on questions of Taiwanese sovereignty (Ko 2026). Such filtering would compress ideological range through hard censorship rather than through alignment training.

The baseline Neutral-response rate is the structural counter-axis to dispersion: it ranges from 81.6% (Qwen3.6 Max Preview) to 6.4% (Gemini 2.5 Flash Lite), a 13× spread largely uncorrelated with compass position. Because our compliance-conditional dispersion is computed only over committed answers, a model that abstains often at baseline has more room to commit under framing; reporting displacement without baseline commitment hides this.

Deployment Risk Profiles: Delegation and Educational Personalization

Controllability is the enabling property of delegated agents. A user who wants an assistant to argue, draft, or triage on their behalf is asking for a system positioned at their values, not the provider’s default. Our dispersion profiles say how reliably each endpoint can be placed there, and how much further it can be pushed. The governance consequence is not obvious, however. When the profile is written by hand, some person authored the normative default and can be held to it. When it is induced from a user’s own interaction history (Jiang et al. 2025) and then tuned by automatic prompt optimization (Yuksekgonul et al. 2025; Khattab et al. 2024), no one explicitly authored it, and the resulting position may be neither inspected nor inspectable by the user it purports to represent. Auditing controllability at the system layer is what makes an induced default legible: it establishes the range within which a learned profile can place a model, and therefore what a delegation is capable of committing its principal to.

The stakes are sharper in personalization contexts where multiple legitimate stakeholders (users, families, institutions, providers, regulators) may disagree about acceptable framing, and where those bounds have to be set normatively, not derived (Kirk et al. 2024). In education the conflict is concrete: personalized content already encodes demographic bias (Weissburg et al. 2025), and unguarded tool design measurably harms learning (Bastani et al. 2025). In child-facing educational systems this becomes a question of authority. Within the range a tutor can be steered across, who sets the default: guardians, schools, providers, or regulators? The operative design question in education is where AI sits in the learning loop, not merely whether it is present (Brcic and Frljic 2026). A steered tutor shapes normative defaults precisely where learners are least equipped to contest them. That is what makes the allocation worth arguing over rather than leaving it to whoever happens to control the instruction layer. What an audit supplies is the range itself, so that the argument runs over a measured space, not an assumed one.

Limitations

This is an observational black-box benchmark; we make no causal claims about training pipelines, RLHF, or internal safety mechanisms. Model endpoints are moving targets, so our results are a dated snapshot (May 2026), not a timeless property of model families. We scope claims to the seven endpoints evaluated; broader patterns are hypotheses for open-source replication, not general assertions about LLMs.

Political Compass as probe, not target. We do not treat Political Compass as a validated ideology theory; we use its items as a fixed, low-dimensional probe for relative response displacement under controlled instruction changes. Our claims are within-instrument – movement, dispersion, asymmetry under matched prompts – not absolute ideological diagnosis. Coordinate noise documented by Röttger et al. (2024a); Ceron et al. (2024) affects absolute placement more than the within-instrument displacement we rely on. The instrument is English-only; paraphrase, instrument-substitution, and multilingual robustness remain future work.

Model scale. Parameter counts are undisclosed for most endpoints, so we cannot test whether dispersion scales with model size as Bernardelle et al. (2025b) report for open-weight families; our tiers are behavioral, not size, groupings. An open-weight replication with known, varied size and a fixed training pipeline would test this directly.

Role-conditioned vs. dialogue-emergent controllability. Our steering operates entirely through fixed system-prompt personas injected before the conversation begins (scope motivated in the Introduction): we measure where a delegated agent can be *placed*, not how it *drifts* once a conversation is under way. Even frontier models track a user’s preferences across a conversation unreliably (Jiang et al. 2025), so whether comparable steerability emerges from dialogue alone, whether induced profiles reach the same positions as hand-written personas, and whether dialogue-emergent controllability shares our saturation ceilings are open questions for future work. The system-prompt layer we probe is also most directly available to platform operators, not end users; induced-profile personalization (§ Introduction) narrows but does not close this gap.

Prompt-time versus training-time steering. We isolate steering at the instruction layer and do not modify weights. Training-time interventions – finetuning, RLHF, activation-level editing (Santurkar et al. 2023; Kim, Evans, and Schein 2025) – move the same behavior through a more permanent channel and can reshape the reachability profile (saturation ceilings, refusal floors) that prompting cannot. The two are complementary; separating their contributions (e.g., whether a saturation ceiling is retraining-removable) requires white-box access and is left to future work.

Forced-choice interface. Our protocol forces a single label per item. This interface may amplify apparent steerability by collapsing ambivalence, hedging, or mixed-position reasoning into discrete directional movement. Our benchmark therefore measures instruction-conditioned response movement under a constrained audit interface, not free-form ideological expression. Open-ended free-form replication is future work.

Bounding within pluralistic-alignment taxonomy. To clarify scope against neighboring concepts, Table 2 maps our measurement against Sorensen’s pluralistic-alignment taxonomy (Sorensen et al. 2024) and the OpenAI Model Spec instruction hierarchy.

Our forced-choice format measures *compliance-conditional dispersion* (how answers shift given that a model answers). A *refusal pilot* (E3) confirms the format

Concept	Measured?	Status
Steerable pluralism	Partially	System-prompt controllability on a political-axis probe
Distributional pluralism	No	Requires population-preference aggregation
Overton-boundary pluralism	No	Requires normative boundary specification
Personalized alignment	No	We measure a prerequisite, not full alignment
Instruction-stack authority	Partially	System layer only; root, developer, and user layers out of scope

Table 2: What this paper does and does not measure.

does not materially suppress refusals. With a softened system prompt that explicitly permitted a REFUSE response on any item, DeepSeek-Chat v3.1 refused on only 1.19% of items across three extreme-framing conditions (0A_3, 0L_3, LO_3; two runs each, 420 question-responses; 5/420: 1/140 under 0A_3 and 4/280 under 0L_3 and LO_3 combined). Refusals remained directionally near-symmetric. Direction preservation held for eight of nine axis×context combinations; the exception was a small economic-axis sign flip under 0L_3. On the model and contexts tested, the forced-choice format does not appear to materially suppress refusal behavior; the E3 pilot is a refusal-floor check on DeepSeek-Chat v3.1, not a full replication across all seven endpoints.

Silent neutral-imputation and the NULL-repair audit. Our original collection pipeline silently imputed Neutral on parse failures. We patched the collector to preserve NULL, re-derived cell scores excluding NULL from denominators, and confirmed via sensitivity analysis that headline-aggregate movement is bounded at $|\Delta| \leq 0.124$ pp on any primary context-axis combination. Repair recovered most failures, and residual missingness is small: in the core data, 64 rows are repair-folded (54 valid repaired responses; 10 persistent missing: 6 refusal-class, 4 unparseable); in the broader repair-backfill audit, 192 candidates produced 159 recoveries and 33 persistent failures. We therefore report complete-case and repair-clean analyses in parallel rather than substituting forced-choice values for refusals.

Temperature ablation (E2). To verify that findings are not artifacts of stochastic sampling, we re-ran the 0A_3 condition for DeepSeek-Chat v3.1 at $T = 0.1$ (near-deterministic) for 2 independent runs. Displacement in (econ, scty) space was 41.2 at $T = 0.1$ versus 45.8 at $T = 0.7$; all 4/4 axes preserved direction. Standard deviation on the economic axis fell from 4.7 ($T = 0.7$) to 3.6 ($T = 0.1$). The effect is temperature-robust; our large context effect sizes ($\eta^2 \approx 0.88\text{--}0.93$ on the primary axes) motivate the main $T = 0.7$ design.

Prompt-length ablation (E1). Higher-intensity framings in our taxonomy are slightly longer than lower-intensity ones, raising a potential prompt-length confound. We tested this directly with a matched-length neutral persona (*The Reflective Inquirer*, 228 words, no ideological content) constructed to match the word-count of our level-3 framings. For DeepSeek-Chat v3.1 the neutral length-matched context produced a mean (econ, scity) L2 displacement of 3.94 points from baseline, versus a mean of 37.78 points for the four partisan extreme conditions, a $9.6\times$ ratio. Prompt length alone cannot account for the steerability signal.

Reachability vs. prompt-family coverage. The compass region occupied by our 12 single-axis contexts (Fig. 2) reflects the prompt families used, not exhaustive model reachability. We therefore ran two 7-model $\times n=5$ pilots whose prompts ask for both axes at once. On the (econ, scity) plane of Fig. 2 (LP_3, LT_3, RP_3, RT_3), no model fills the square: per-model coverage is 10.6–32.9% of the plane, and the culturally bundled Left-Progress $\leftrightarrow$ Right-Tradition diagonal is reached 5.8–12.9 pp closer than the off-diagonal corners on all seven models. A second pilot on the (econ, govt) plane (LA_3, LL_3, RA_3, RL_3) repeats the pattern: coverage 12.8–38.7%, and a Left-Lib $\leftrightarrow$ Right-Auth diagonal advantage of 2.7–11.4 pp on six of seven models, with Grok-4.3 near-symmetric on that plane only. This bounded coverage is consistent both with model-side compression of off-diagonal combinations and with axis coupling in the 8values items themselves; separating the two would require an instrument-substitution study. Both pilots are reported in full in the supplementary appendix.

Open-ended reasoning, multilingual transfer, and open-source baselines (Bernardelle et al. 2025b) are out of scope.

Reproducibility

We release everything needed to repeat the study: the 12 persona prompts and baseline prompt, the 70-item instrument, per-question axis weights, raw model answers ($N = 63,700$), and analysis scripts for the variance decomposition, bootstrap intervals, signed-axis projection, and cross-model permutation test. Running the package rebuilds every processed file from the raw answers and re-checks every reported number. It is archived at Zenodo (DOI 10.5281/zenodo.21489805) (Bućan et al. 2026) and browsable at <https://github.com/mbricic/llm-political-steerability>; § Models and Inference Settings lists the exact API identifiers so later versions can be benchmarked under the same protocol.

Conclusion

The question “Is this model politically biased?” is less useful than it appears. A more productive question, and the one our benchmark operationalizes, is whether value-sensitive outputs are instruction-conditionable in predictable, asymmetric, and auditable ways. Within our forced-choice probe, the answer across seven leading commercial endpoints is yes. Context framing accounts for roughly 88% of economic-axis variance and 93% of society-axis variance; inter-model baseline differences account for under 3%. A further 5–8% sits in the context $\times$ model interaction. Models do

differ, then, but principally in how they respond to framing, not in where they start. Controllability is non-uniform across dispersion/symmetry/saturation/refusal-floor dimensions, and per-question shifts converge across models at the level of observed outputs (a large majority of PC items receive unanimous-sign agreement under authoritarian framing) without our being able to adjudicate the underlying system-layer driver.

We recommend that future political evaluations report controllability profiles alongside position estimates, decompose directional claims into displacement and proximity, treat reliability concentration as part of the empirical result, and audit the instruction-bearing layers that shape deployed behavior. The question is not only whether a model has a political location, but *who can move it, how far, under what authority, through which instruction layer, and with what asymmetries*.

Adverse Impact and Positionality

Adverse impact. This dispersion-first framing is dual-use: showing that frontier LLMs comply with directional framings, including authoritarian ones, could inform red-teaming or adversarial influence. We report aggregate magnitudes, not per-question persuasion recipes, and release no optimisation pipeline for influence. The deeper risk is deployment: a system steerable anywhere on a value axis concentrates normative authority in the instruction layer, whether set by a platform or induced from a user profile (§Deployment Risk Profiles), which argues for disclosing steering ranges alongside baselines. Our framings come from a public political-compass tradition, not purpose-built persuasion datasets. We judge the benchmark’s safety value to outweigh its capability uplift, though this judgement is debatable.

Positionality. The authors are academic researchers in computer science and AI ethics at a European university, with no commercial relationship to the seven providers evaluated. We bring our own priors: treating “models that match the user’s framing” as worth measuring reflects a position on AI behavior and pluralism. Our metrics are symmetric across directions, so readers with different priors can still interrogate the data.

Acknowledgments

This research was funded by the European Union NextGenerationEU through the National Recovery and Resilience Plan 2021–2026, under the institutional grant of the University of Zagreb Faculty of Electrical Engineering and Computing, project Value-aligned and interpretable optimization and reasoning (VALOR).

Declaration on the Use of AI Assistance

Generative AI assistants were used only for language polishing of author-written text and for assistance with analysis and experiment code. All research design, analyses, interpretations, and claims are the authors’ own; every AI-assisted output was reviewed and verified by the authors, who take full responsibility for the content of this paper.

References

- Adams, J.; Hu, B.; Veenhuis, E.; Joy, D.; Ravichandran, B.; Bray, A.; Hoogs, A.; and Basharat, A. 2025. Steerable Pluralism: Pluralistic Alignment via Few-Shot Comparative Regression. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 8(1): 15–25.
- Aldahoul, N.; et al. 2025. Large Language Models are often politically extreme, usually ideologically inconsistent, and persuasive even in informational contexts. *arXiv preprint arXiv:2505.04171*.
- Anthropic. 2025. Measuring political bias in Claude. Anthropic research report.
- Bastani, H.; Bastani, O.; Sungu, A.; Ge, H.; Kabakçı, Ö.; and Mariman, R. 2025. Generative AI Without Guardrails Can Harm Learning: Evidence from High School Mathematics. *Proceedings of the National Academy of Sciences*, 122: e2422633122. Correction: DOI 10.1073/pnas.2518204122.
- Batzner, J.; Stocker, V.; Schmid, S.; and Kasneci, G. 2025. GermanPartiesQA: Benchmarking Commercial Large Language Models and AI Companions for Political Alignment and Sycophancy. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 8(1): 330–342.
- Bernardelle, P.; et al. 2025a. Mapping and Influencing the Political Ideology of Large Language Models using Synthetic Personas. In *Companion Proceedings of the ACM Web Conference 2025 (WWW '25 Companion)*.
- Bernardelle, P.; et al. 2025b. Political Ideology Shifts in Large Language Models. *arXiv preprint arXiv:2508.16013*.
- Brcic, M.; and Frlić, S. 2026. The Effortless Trap: Productive Struggle, AI, and the Illusion of Learning. *arXiv preprint arXiv:2606.26181*.
- Brcic, M.; and Yampolskiy, R. V. 2023. Impossibility Results in AI: A Survey. *ACM Computing Surveys*, 56(1): 1–24.
- Bučan, B.; Sočec, N.; Isufi, S.; Granić, M.; Hobor, L.; Krajna, A.; Kovac, M.; and Brcic, M. 2026. Auditing Alignment Controllability in LLMs via Political Axes: Reproducibility Package (Code and Data). Concept DOI, all versions.
- Ceron, T.; Falk, N.; Barić, A.; Nikolaev, D.; and Padó, S. 2024. Beyond Prompt Brittleness: Evaluating the Reliability and Consistency of Political Worldviews in LLMs. *Transactions of the Association for Computational Linguistics*, 12: 1378–1400.
- Chang, T.; Schnabel, T.; Swaminathan, A.; and Wiens, J. 2025. A Course Correction in Steerability Evaluation: Revealing Miscalibration and Side Effects in LLMs. *arXiv preprint arXiv:2505.23816*.
- Cintas, C.; Rateike, M.; Miehl, E.; Daly, E.; and Speakman, S. 2025. Localizing Persona Representations in LLMs. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 8(1): 630–642.
- Cui, J.; Chiang, W.-L.; Stoica, I.; and Hsieh, C.-J. 2025. OR-Bench: An Over-Refusal Benchmark for Large Language Models. In Singh, A.; Fazel, M.; Hsu, D.; Lacoste-Julien, S.; Berkenkamp, F.; Maharaj, T.; Wagstaff, K.; and Zhu, J., eds., *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, 11515–11542. PMLR.
- Hackenburg, K.; and Margetts, H. 2024. Evaluating the Persuasive Influence of Political Microtargeting with Large Language Models. *Proceedings of the National Academy of Sciences*, 121(24): e2403116121.
- Hackenburg, K.; Tappin, B. M.; Hewitt, L.; Saunders, E.; Black, S.; Lin, H.; Fitt, C.; Margetts, H.; Rand, D. G.; and Summerfield, C. 2025. The levers of political persuasion with conversational artificial intelligence. *Science*, 390(6777): eaea3884.
- Hartmann, J.; Schwenzow, J.; and Witte, M. 2023. The political ideology of conversational AI: Converging evidence on ChatGPT’s pro-environmental, left-libertarian orientation. *arXiv preprint arXiv:2301.01768*.
- Hu, B.; Ray, B.; Leung, A.; Summerville, A.; Joy, D.; Funk, C.; and Basharat, A. 2024. Language Models are Alignable Decision-Makers: Dataset and Application to the Medical Triage Domain. *arXiv:2406.06435*.
- Jakesch, M.; et al. 2023. Co-Writing with Opinionated Language Models Affects Users’ Views. In *Proceedings of CHI*. ArXiv:2302.00560.
- Jiang, B.; Hao, Z.; Cho, Y.-M.; Li, B.; Yuan, Y.; Chen, S.; Ungar, L.; Taylor, C. J.; and Roth, D. 2025. Know Me, Respond to Me: Benchmarking LLMs for Dynamic User Profiling and Personalized Responses at Scale. *arXiv:2504.14225*.
- Kabir, S. 2025. When Models Refuse: Political Steerability and Feature Richness as Measures of Ideological Depth. *arXiv preprint arXiv:2508.21448*.
- Khattab, O.; Singhvi, A.; Maheshwari, P.; Zhang, Z.; Santhanam, K.; Vardhamanan, S.; Haq, S.; Sharma, A.; Joshi, T. T.; Moazam, H.; Miller, H.; Zaharia, M.; and Potts, C. 2024. DSPy: Compiling Declarative Language Model Calls into State-of-the-Art Pipelines. In *International Conference on Learning Representations (ICLR)*.
- Kim, J.; Evans, J.; and Schein, A. 2025. Linear Representations of Political Perspective Emerge in Large Language Models. In *International Conference on Learning Representations (ICLR)*. Oral; *arXiv:2503.02080*.
- Kirk, H. R.; et al. 2024. The Benefits, Risks and Bounds of Personalizing the Alignment of Large Language Models to Individuals. *Nature Machine Intelligence*, 6(4): 383–392.
- Ko, J.-C. 2026. Bilingual Bias in Large Language Models: A Taiwan Sovereignty Benchmark Study. *arXiv preprint arXiv:2602.06371*.
- Li, J.; Peris, C.; Mehrabi, N.; Goyal, P.; Chang, K.-W.; Galstyan, A.; Zemel, R.; and Gupta, R. 2024. The Steerability of Large Language Models Toward Data-Driven Personas. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 7290–7305. ArXiv:2311.04978.
- Miehl, E.; et al. 2025. Evaluating the Prompt Steerability of Large Language Models. In *Proceedings of NAACL. ACL Anthology 2025.naacl-long*.400.

- OpenAI. 2025. Defining and Evaluating Political Bias in LLMs. OpenAI blog post.
- Peng, T.-Q.; Yang, K.; Lee, S.; Li, H.; Chu, Y.; Lin, Y.; and Liu, H. 2026. Beyond partisan leaning: a comparative analysis of political bias in large language models. *Journal of Information Technology & Politics*. ArXiv:2412.16746.
- Perez, E.; et al. 2023. Discovering Language Model Behaviors with Model-Written Evaluations. In *Findings of the Association for Computational Linguistics (ACL Findings)*. ArXiv:2212.09251.
- Potter, Y.; et al. 2024. Hidden Persuaders: LLMs’ Political Leaning and Their Influence on Voters. In *Proceedings of EMNLP*, 4244–4275.
- Röttger, P.; et al. 2024a. Political Compass or Spinning Arrow? Towards More Meaningful Evaluations for Values and Opinions in Large Language Models. In *Proceedings of ACL*, 15295–15311.
- Röttger, P.; et al. 2024b. XSTest: A Test Suite for Identifying Exaggerated Safety Behaviours in Large Language Models. In *Proceedings of NAACL*, 5377–5400.
- Rozado, D. 2023. The Political Biases of ChatGPT. *Social Sciences*, 12(3): 148.
- Rozado, D. 2024. The Political Preferences of LLMs. *PLOS ONE*, 19(7): e0306621.
- Sakhawat, A.; et al. 2026. Political Alignment in Large Language Models: A Multidimensional Audit of Psychometric Identity and Behavioral Bias. *arXiv preprint arXiv:2601.06194*.
- Santurkar, S.; et al. 2023. Whose Opinions Do Language Models Reflect? In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Sharma, M.; et al. 2024. Towards Understanding Sycophancy in Language Models. In *International Conference on Learning Representations (ICLR)*.
- Smith-Vaniz, N.; Lyon, H.; Steigner, L.; Armstrong, B.; and Mattei, N. 2025. Investigating Political and Demographic Associations in Large Language Models Through Moral Foundations Theory. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 8(3): 2419–2430.
- Sorensen, T.; et al. 2024. Position: A Roadmap to Pluralistic Alignment. In *International Conference on Machine Learning (ICML)*. ArXiv:2402.05070.
- Tseng, Y.-M.; Huang, Y.-C.; Hsiao, T.-Y.; Chen, W.-L.; Huang, C.-W.; Meng, Y.; and Chen, Y.-N. 2024. Two Tales of Persona in LLMs: A Survey of Role-Playing and Personalization. arXiv:2406.01171.
- Weissburg, I.; Anand, S.; Levy, S.; and Jeong, H. 2025. LLMs are biased teachers: Evaluating llm bias in personalized education. In *Findings of the Association for Computational Linguistics: NAACL 2025*, 5650–5698.
- Yuksekgonul, M.; Bianchi, F.; Boen, J.; Liu, S.; Lu, P.; Huang, Z.; Guestrin, C.; and Zou, J. 2025. Optimizing generative AI by backpropagating language model feedback. *Nature*, 639(8055): 609–616.
- Zhang, Z.; Rossi, R. A.; Kveton, B.; Shao, Y.; Yang, D.; Zamani, H.; Derroncourt, F.; Barrow, J.; Yu, T.; Kim, S.; Zhang, R.; Gu, J.; Derr, T.; Chen, H.; Wu, J.; Chen, X.; Wang, Z.; Mitra, S.; Lipka, N.; Ahmed, N.; and Wang, Y. 2024. Personalization of Large Language Models: A Survey. *arXiv preprint arXiv:2411.00027*.