People Analytics Network Census (PANC) 2026: Identifying influencers in the global People Analytics community

The PANC Core Team*

2026

Abstract

The People Analytics Network Census (PANC) maps the global professional network of People Analytics (PA) practitioners and tracks how it evolves over time. One of its objectives is to identify the different kinds of influencers in the community. This paper documents the methodology used in the 2026 wave to identify those influencers and records the empirical basis for each parameter choice. The analysis produces several lists. A weighted PageRank over three name-generator networks yields a Top 30 Overall list, a Top 3-per-country list, and a Rising Stars list of earlier-career people already drawing influence. A marker-adjusted observed-versus-expected score, which corrects for variation in how freely individual nominators attribute each quality, yields four category lists (Vision and Ideas, Mentorship, Technical Expertise, Networking), each split between internal-team practitioners and everyone else. Influence is computed over the full nominated network so that everyone's nominations count; consent to be named (and, for some lists, a minimum level of corroborating nominations) then governs who appears in each published list. The survey instrument is reproduced in full in the appendix.

1 Background: the People Analytics Network Census

The People Analytics Network Census (PANC) is a community initiative that aims to map and understand the global network of People Analytics professionals, and to understand how that network changes over time. It is run by a six-person core team drawn from across the field: Patrick Coolen (KennedyFitch), Matthew Diabes (NYU), Stephanie Murphy (Society for People Analytics), Maria Nolazco Masson (IPSY), Andrew Pitts (Polinode) and Richard Rosenow (Ikona Analytics). The census is distributed in partnership with Directionally Correct, Insight222, People Analytics World and the Society for People Analytics.

The census has three aims. The first is to map the global PA professional network. The second is to identify the different types of influencer within it. The third is to deliver insights back to the community and to individual participants. Three deliverables follow from these aims. Personal Network Reports give each respondent a view of where they sit in the overall network and how they compare to others, and name no one other than the recipient. A community webinar presents the core analysis and its main findings. An NYU research component, led by Matthew Diabes, uses a short battery of validated IO-psychology measures included in the survey to study the relationship between certain individual differences and network position; it is covered by NYU IRB approval (IRB-FY2025-9279).

*Patrick Coolen (KennedyFitch), Matthew Diabes (NYU), Stephanie Murphy (Society for People Analytics), Maria Nolazco Masson (IPSY), Andrew Pitts (Polinode), Richard Rosenow (Ikona Analytics).

1.1 Survey design and data

The survey is administered through Polinode's Active Organizational Network Analysis (ONA) functionality. It is a name-generator instrument: each respondent nominates the people who matter to them across three relationship contexts, which makes the resulting network directed (respondent to nominee). The instrument has four sections, covering informed consent, the network name-generators, demographic and background questions, and the optional NYU research battery. All four are reproduced in Appendix A.

The three name-generator questions each ask for up to ten nominations (Appendix A, Q4–Q6). Q4 asks for key working relationships within the PA team at the respondent's current workplace. Q5 asks for maintained relationships from a former workplace's PA team. Q6 asks for important relationships from the wider PA community that the respondent does not currently work with and has not worked with in the past, such as relationships formed through conferences, professional associations, academia, social media or vendors.

A fourth question, Q7, qualifies each nominated tie. For every person named in Q4–Q6, the respondent indicates whether that person is, for them, a source of Vision and Ideas, Mentorship, Technical Expertise or Networking. Any combination may be selected, including none. These four tie types are the basis for the category-influencer lists.

All responses are anonymised before analysis. Each person is keyed by an opaque identifier, with names removed, and only the core team has access to the individual-level anonymised data. The 2026 analysis works entirely from anonymised identifiers. Mapping a result back to a named individual requires a separately held identifier-to-name mapping, and any public naming is governed by the respondent's consent (Q2).

2 Identifying influencers in 2026

The 2026 wave produces its influencer lists from two methods. The overall method is a weighted PageRank over the network structure, and from it we derive an overall list, a per-country list, and a list of earlier-career rising stars. The category method is a per-tie score that rewards being sought out for a specific quality, reported separately for internal-team practitioners and for everyone else. We describe each in turn below.

PageRank and the category scores are computed over the *full* nominated network, so that everyone's nominations count toward others' standing. Whether a person then appears in a given published list is governed by consent and, for some lists, by corroboration. Only people who *consented to being named* as an influencer (question Q2) appear in any published list; a person who declined to be named is never published, however high they rank. Some lists additionally require a minimum level of corroborating nominations, as noted where they are defined.

2.1 Overall influence: weighted PageRank

Overall influence is computed with PageRank over the union of all three name-generator networks (Q4, Q5 and Q6). The three networks are not treated identically. Each is given a different edge weight, on the view that standing in the wider community is the strongest signal of community-level influence, whereas a current-workplace tie reflects organisational structure as much as influence. The weights are set out in Table 1.

PageRank is run on this single weighted directed graph. A person scores highly when many people, and in particular well-connected people, nominate them as an important community relationship. The weights are a deliberate starting point and are straightforward to revise.

Network	Survey question	Edge weight
Wider community	Q6	0.50 (highest)
Former workplace	Q5	0.30
Current workplace	Q4	0.20 (lowest)

Table 1: Edge weights applied to the three name-generator networks in the weighted-PageRank computation of overall influence.

The published Top 30 Overall Influencers list comprises the thirty highest-scoring eligible people.

From the same ranking we derive a Top 3 Influencers by Country list. For every country, we take the three highest-scoring people who clear a minimum corroboration threshold of at least two incoming nominations across the three networks, so that a country is not represented by someone named only once. Unlike a previous outside-the-United-States formulation, this list includes the United States and does not exclude people who already appear in the Overall list; it is simply the strongest, sufficiently-corroborated people in each country.

2.2 Rising stars

The same overall ranking also supports a Rising Stars list, which surfaces people who are relatively early in their careers yet already draw influence in the network. A person qualifies for this list when they are earlier-career on two measures at once — no more than fifteen years of total work experience (Q20) *and* no more than eight years of People Analytics experience (Q21) — and when their standing is corroborated by at least two incoming nominations. The two experience ceilings sit at or just below the community medians (about seventeen years of total experience and nine years of People Analytics experience in this wave), so the list reflects genuinely earlier-career people. People already in the Top 30 Overall list are excluded, since the intent is to highlight emerging influence rather than to restate the established leaders. The list reports the ten highest-ranked qualifying people.

2.3 Category influence: a marker-adjusted score

For each of the four tie types (Vision and Ideas, Mentorship, Technical Expertise, Networking) the goal is to find the people the community most associates with that quality. The intuitive measure is the proportion of a person's nominations that carry the tie, but this simple proportion has a specific weakness.

The problem. Nominators differ greatly in how freely they attribute a given quality. In the 2026 data, the per-nominator rate of tagging Technical Expertise spans the entire range from 0 to 1: some respondents tag no one as technical, while others tag everyone they named. Global tag rates also vary by dimension, from about 0.27 for Mentorship to about 0.60 for Vision and Ideas. A plain proportion is therefore unfair. It under-credits a person whose endorsers happen to be demanding markers, and over-credits a person whose endorsers are lenient, as an accident of who named them.

The adjustment. We replace the proportion with a shrunk observed-versus-expected ratio. For a given tie type, the calculation proceeds in three steps. First, each nominator s is assigned a base rate r_s, defined as how often they attribute this quality to anyone they nominate, smoothed toward the global rate so that a nominator with very few nominations is not judged as an extreme case.

Second, for each target the expected number of ties is $\exp = \sum_s r_s$, summed over the people who named them; this is the number a typical person nominated by that same set of nominators would receive. Third, the score is

$$\text{score} \;=\; \frac{\text{obs} + K}{\exp + K}, \tag{1}$$

where obs is the number of ties actually received and K is a shrinkage constant. A score above 1 means the person was sought out for this quality more than expected, given who nominated them and how demanding those nominators are.

This score has two properties the project requires. It rewards endorsement by selective markers, because a technical tie from someone who rarely calls anyone technical contributes more than one from someone who tags everyone. It also does not reward volume, because the denominator scales with the number of nominators; being named by many people does not by itself inflate the score. Running PageRank on a single tie network would not have this second property, as it would re-introduce a popularity bias. People are ranked within each category subject to a minimum number of total nominations, discussed in Section 2.4.

For each of the four categories we report two short lists side by side: the strongest internal-team practitioners and the strongest people from everyone else. The split is made on the respondent's primary-role answer (Q14): one list is the people whose role is an internal People Analytics team, the other is everyone with a different role. Reporting the two separately keeps a practitioner who is sought out for, say, technical expertise from being measured against vendors, consultants and academics whose roles position them differently, and makes each list more directly useful to its audience.

2.4 Parameter choices and their justification

Parameters were chosen empirically rather than by assertion. The main tool is split-half reliability. We randomly split the nominations (the network edges) into two halves, compute the ranking on each half independently, and measure how well the two rankings agree, using rank correlation and top-20 overlap averaged over many random splits. A good parameter value is one whose ranking is stable under this resampling of the observed nominations. Because both halves are drawn from the same survey wave and rank the same set of people, this measures internal consistency under edge-level resampling rather than validation on a genuinely independent sample; the values should be read as an upper bound on cross-sample reliability. It is nonetheless well suited to its purpose here, which is to choose parameters by comparing reliability across candidate values: any shared inflation from the network's latent structure affects all candidates similarly, so the comparison, and hence the selected value, is robust.

Shrinkage $K = 5$. The constant K in Equation (1) adds K neutral pseudo-observations sitting at expectation, which pulls thinly-evidenced scores toward the no-signal value of 1, and pulls harder the smaller the sample. Across all four dimensions, $K = 0$ (no shrinkage) is the least reliable setting. Reliability rises with K and reaches a plateau around $K \approx 5$ to 8, most visibly on the thin-signal Vision and Ideas dimension. We use $K = 5$. It sits at the reliability plateau, works as a single value across all four categories, and increasing it further only shrinks genuine signal for no gain. A method-of-moments estimate of the ideal shrinkage was too numerically unstable on data of this size to identify a value, which is why the reliability curve drives the choice.

Minimum nominations of 8. A floor on total nominations is needed because the expected denominator is unstable for people with very few nominators. Since the shrinkage already controls small samples, the floor can sit at eight total nominations without admitting artifacts. People who enter the lists at eight or nine nominations show genuine lift, and cases that would otherwise reach 100 per cent on a tiny denominator are pulled to the middle of the list rather than the top. A single floor is applied across all four dimensions; the differing base rates are absorbed by the shrinkage rather than requiring separate thresholds per dimension.

An extension we tested and rejected. One natural refinement is to weight a nominator's judgement by their own standing in that dimension, on the theory that a recognised technical expert is a better judge of technical expertise than a non-specialist. This is the recursive, eigenvector-style idea that underlies PageRank, and we tested a simple one-step version of it. Under the same split-half reliability test, leaning on expert judges made the rankings monotonically less reliable in every dimension. The reason is structural: expert judges are a small subset, and each target receives only a few of their votes, so up-weighting them discards the stabilising volume of ordinary nominators and concentrates the score on a handful of noisy votes. We therefore do not weight judges by competence. The harshness correction described in Section 2.3, by comparison, passes a construct-validity check: the people it moves are those with the most skewed nominator pools, it moves them in the correct direction (harsh-endorsed people up, lenient-endorsed people down), and it leaves people with average nominator pools close to where they started.

3 Interpretation and scope

Four points help place the results in context. First, the category lists are best read as groups of strong candidates rather than an exact ranking: each captures the people the community clearly associates with that quality, while the precise order within a list carries less weight. The same reading applies to the per-country and rising-stars lists, which identify strong, sufficiently-endorsed people in a given setting rather than asserting a fine-grained order. Second, the observed-versus-expected score is built on the principle that a tie from a more selective nominator is more informative, which the construct-validity checks in Section 2.4 support; it measures standing as expressed through the community's own nominations rather than any external benchmark of expertise. Third, the overall-influence weights of 0.50, 0.30 and 0.20 are a deliberate, transparent starting point that performs well in practice and can be refined as the analysis matures across future waves. Fourth, the lists are deliberately shaped to be fair and useful as well as accurate: requiring a minimum level of corroborating nominations keeps thinly-evidenced names off the per-country and rising-stars lists, and separating practitioners from other roles compares like with like.

A Survey instrument (PANC 2026)

The questions as administered are reproduced below. Response scales are noted where they bear on the analysis. Sections 1 to 3 are shown in full. Section 4, the optional NYU research battery (Q28–Q31), is summarised, as it does not feed the influencer analysis.

Section 1: introduction and informed consent

Q1. Do you consent to the collection of your data for the purpose of the People Analytics Network Census? *(consent)*

Q2. If you are identified as an influencer, do you consent to being named publicly as an influencer? You may proceed even if you answer no. *(consent to be named)*

Q3. Would you like to receive a copy of a personal network report based on your answers and the anonymized answers of others? *(report opt-in)*

Section 2: your connections to the PA community

General instructions: nominees are selected by name, and new names may be added with the organisation given in brackets. Each of Q4–Q6 is capped at ten nominations, and a full ten is not required.

Q4. Current workplace. Nominate up to ten people who are your key working relationships within the PA team at your current workplace.

Q5. Former workplace. Nominate up to ten people who were key working relationships within the PA team at a former workplace and with whom you have maintained contact.

Q6. Wider community. Nominate up to ten important relationships from the PA community that you do not currently work with and have not worked with in the past.

Q7. Tie qualification (per nominee). For each person nominated in Q4–Q6, indicate whether they are, for you, a source of: Vision and Ideas (strategic insight or innovative thinking); Mentorship (mentorship or career advice); Technical Expertise (help on difficult technical problems); or Networking (thoughtful introductions to others). None, one, or several may be selected.

Section 3: about you and your background

All questions in this section are optional.

Q8. Age group *(single-select bands)*.

Q9. Gender identity *(single-select)*.

Q10. Race or ethnicity *(US respondents only; single-select)*.

Q11. Country of residence *(free text)*.

Q12. City of residence *(free text)*.

Q13. Current job title *(free text)*.

Q14. Which best describes your primary role? People Analytics internal team; HR or People; PA vendor; Consultant; Academic; Student; or Other *(single-select; drives skip logic for Q16–Q18)*.

Q15. Current seniority level *(Individual Contributor through C-suite)*.

Q16. PA team scope of responsibility *(multi-select; internal PA teams only)*.

Q17. PA team size *(single-select; internal PA teams only)*.

Q18. PA team maturity *(Initial, Limited, Systematic or Strategic; internal PA teams only)*.

Q19. Approximate organisation size *(single-select bands)*.

Q20. Total years of full-time work experience *(numeric)*.

Q21. Years of experience in People Analytics *(numeric)*.

Q22. Have you taken a break of three or more months for maternity, paternity, or carer's leave? *(Yes or No)*

Q23. Time since returning to work *(numeric years; conditional on Q22 = Yes).*

Q24. Current career orientation *(single-select).*

Q25. Have you managed people in previous roles? *(Yes or No)*

Q26. Have you managed managers in previous roles? *(Yes or No)*

Q27. Community connectedness: seven statements covering local, national and global PA connectedness, learning from and contributing to the community, information access, and intent to remain in PA, each rated on a five-point agree-disagree scale.

Section 4: NYU research questions (optional)

A separate optional battery of validated IO-psychology measures contributed to the NYU research component (NYU IRB-FY2025-9279). It comprises four matrix questions: Q28, well-being and psychological richness (six items, seven-point scale); Q29, self-monitoring (four items, six-point scale); Q30, brokerage in the *tertius iungens* or connecting sense (six items, seven-point scale); and Q31, brokerage in the *tertius gaudens* or separating sense (six items, seven-point scale). These items do not feed the influencer analysis described in this paper.