


What Do NLP Researchers Believe?

Results of the NLP Community Metasurvey

Julian Michael,¹ Ari Holtzman,² Alicia Parrish,^{3*} Aaron Mueller,⁴ Alex Wang,^{5*}
Angelica Chen,¹ Divyam Madaan,¹ Nikita Nangia,¹
Richard Yuanzhe Pang,¹ Jason Phang,¹
Samuel R. Bowman^{1,6*}

<https://nlpsurvey.net/results>



*work done while at NYU

Analyze the results yourself on our dashboard!

Motivation

- Developments in NLP are **hard to predict & becoming higher stakes**.
- This raises **a lot of questions** where there is **a lot of disagreement**.
- It's hard to know via social media what most people **actually think**.
- This makes it **harder to communicate and build shared agendas**.
- Can lead researchers to avoid promising topics, waste effort, and argue on the basis of assumptions not held by their audience.
- Worst case scenario: **a fractured discourse**
- We can do something about this — **run a survey!**

Methodology

- Each question is an opinion with a 4-point Likert scale plus 3 options:

Q1-1. Private firms have too much influence in guiding the trajectory of the field.


- Agree
- Weakly agree
- Weakly disagree
- Disagree
- Other:
 - Insufficiently informed on the issue
 - Question is ill-posed
 - Prefer not to say

- Then respondents *predict what others will say*, and we can compare!

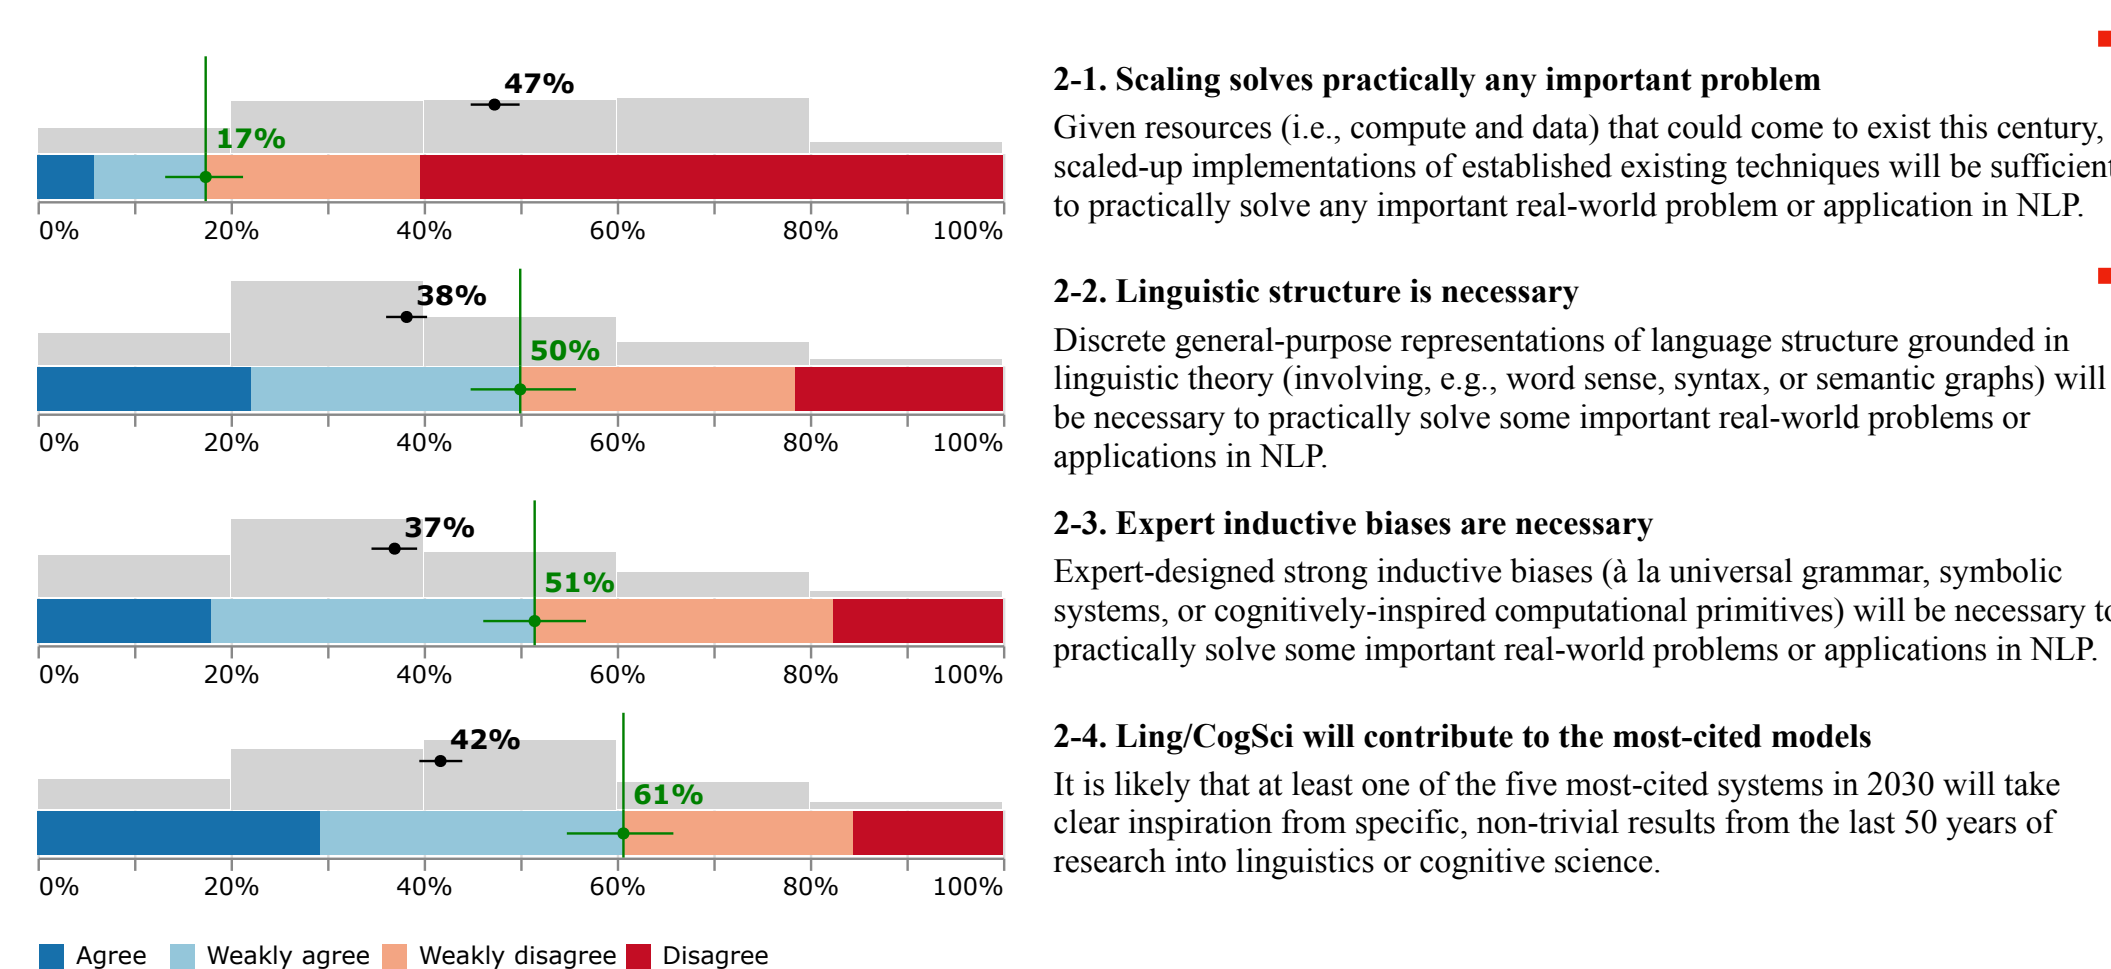
Of those on the agree/disagree spectrum, what percentage of community members do you think will mark "agree" or "weakly agree"?

0-20% 20-40% 40-60% 60-80% 80-100%

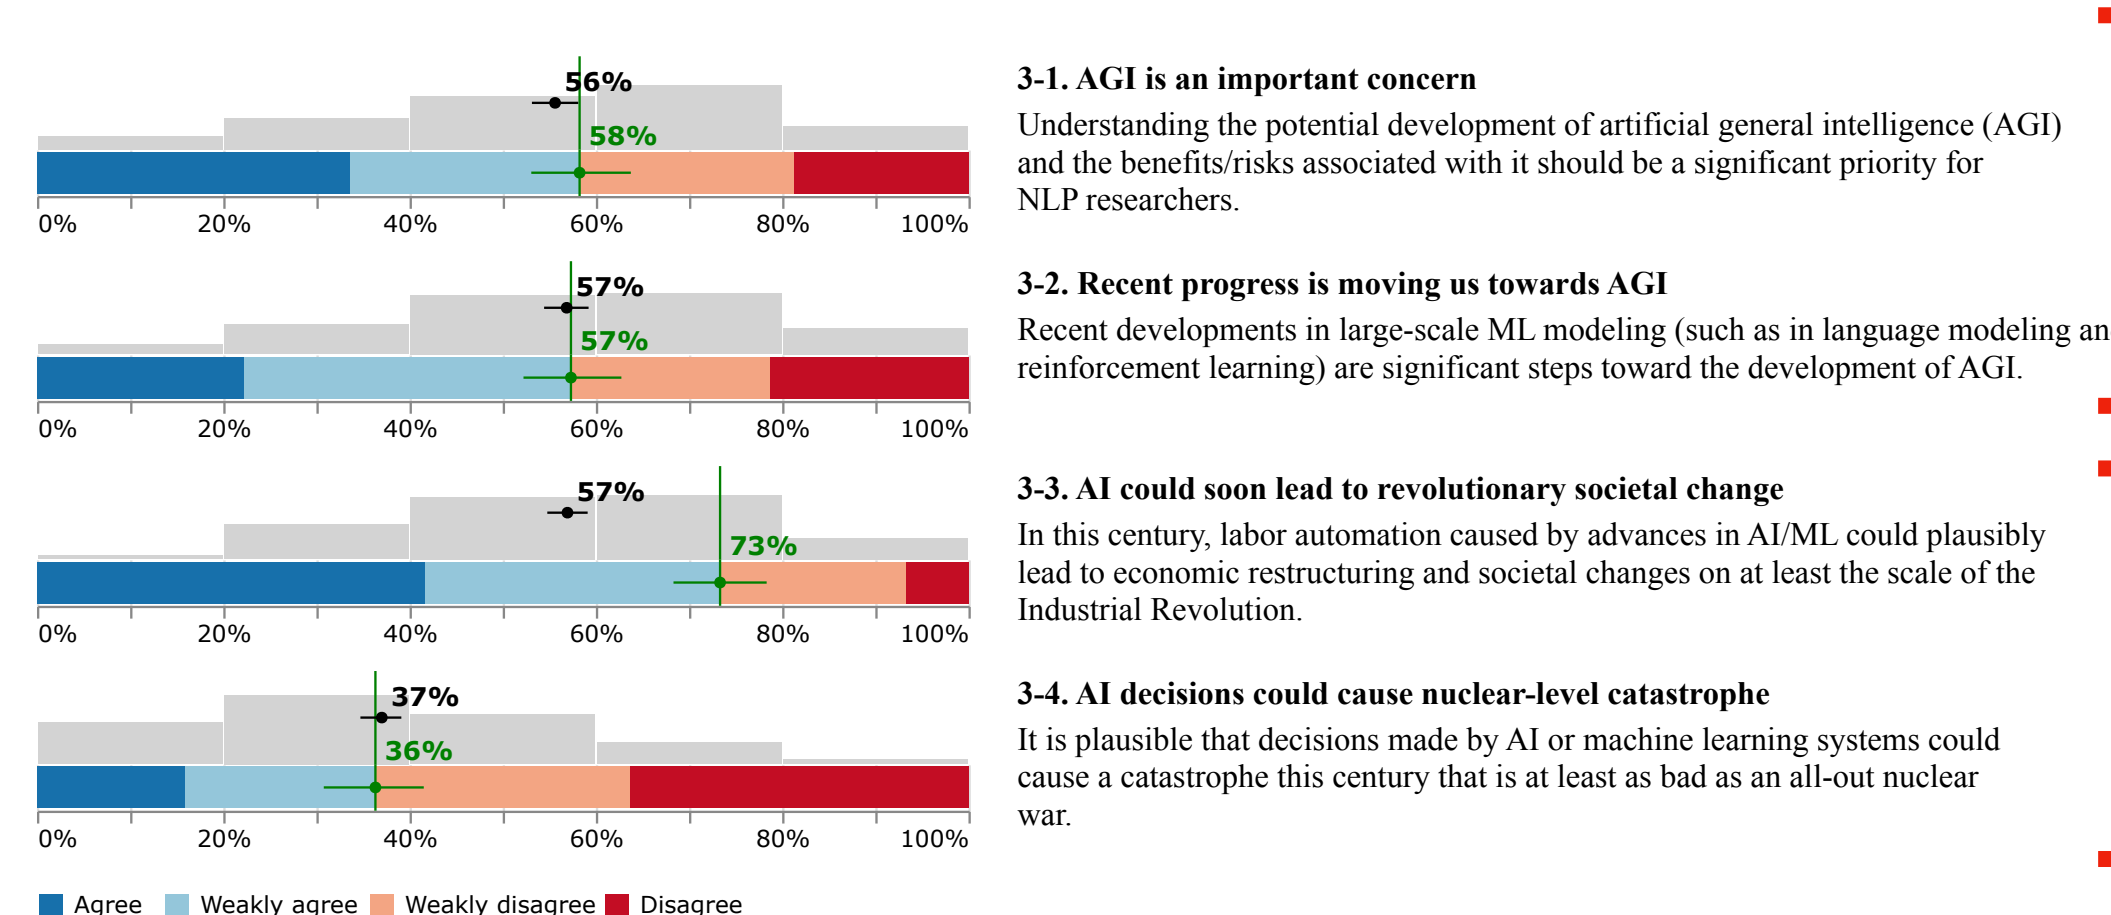
1. State of the Field



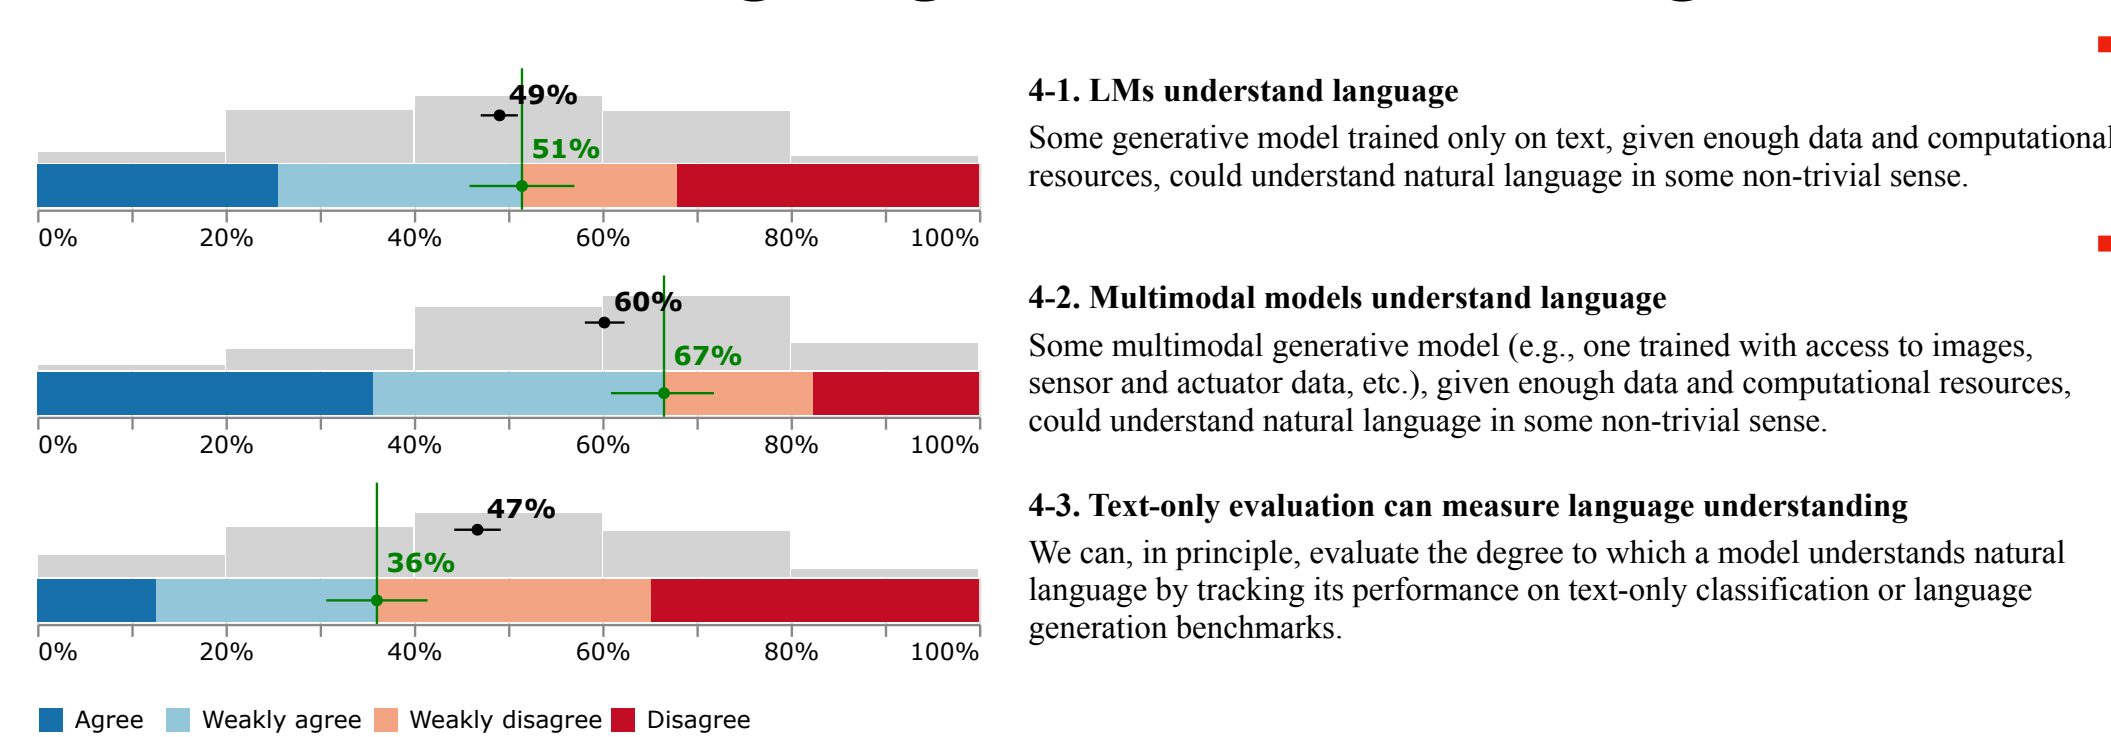
2. Scale, Inductive Bias, and Adjacent Fields



3. AGI and Major Risks



4. Language Understanding



Observations

Respondents skew senior, academic, and online

A plurality (41%) of respondents are senior, 73% work in academia, and 89% publish for their job. Most (76%) use Twitter.

The USA is overrepresented

The USA is overrepresented (58%, compared to 35% of ACL members), and Asia underrepresented but the gender ratio tracks available ACL statistics.

Industry is seen as influential

Industry is seen as highly influential — 77% think it has too much influence, and 86% think it will produce the most widely-cited research. The prevalence of these views are underestimated by 15–20%.

NLP seen as dubious science

67% believe that a majority of NLP publications have dubious scientific value.

People want to prioritize interdisciplinary insights

82% believe we should incorporate more insights from adjacent fields / domain sciences — many more than the 53% predicted.

Scaling maximalism is rare

Only 17% believe in *scaling maximalism* — the view that scale will solve all practical problems. The view's prevalence is greatly overestimated, at 47%.

Problem formulation and task design: a frontier?

The most influential advances of the next 10 years are expected to be in *problem formulation and task design* rather than scale, hardware, data, or algorithms.

AGI and "understanding" are controversial still

Whether AGI is an important concern, and whether LMs understand language, are split down the middle.

AI is seen as plausibly having massive impact on society

73% think AI could plausibly bring revolutionary societal change, and 36% think catastrophic risk from AI is plausible, even though a vast majority (87–89%) thinks NLP is net good for the world.

There is concern for carbon but skepticism of regulation

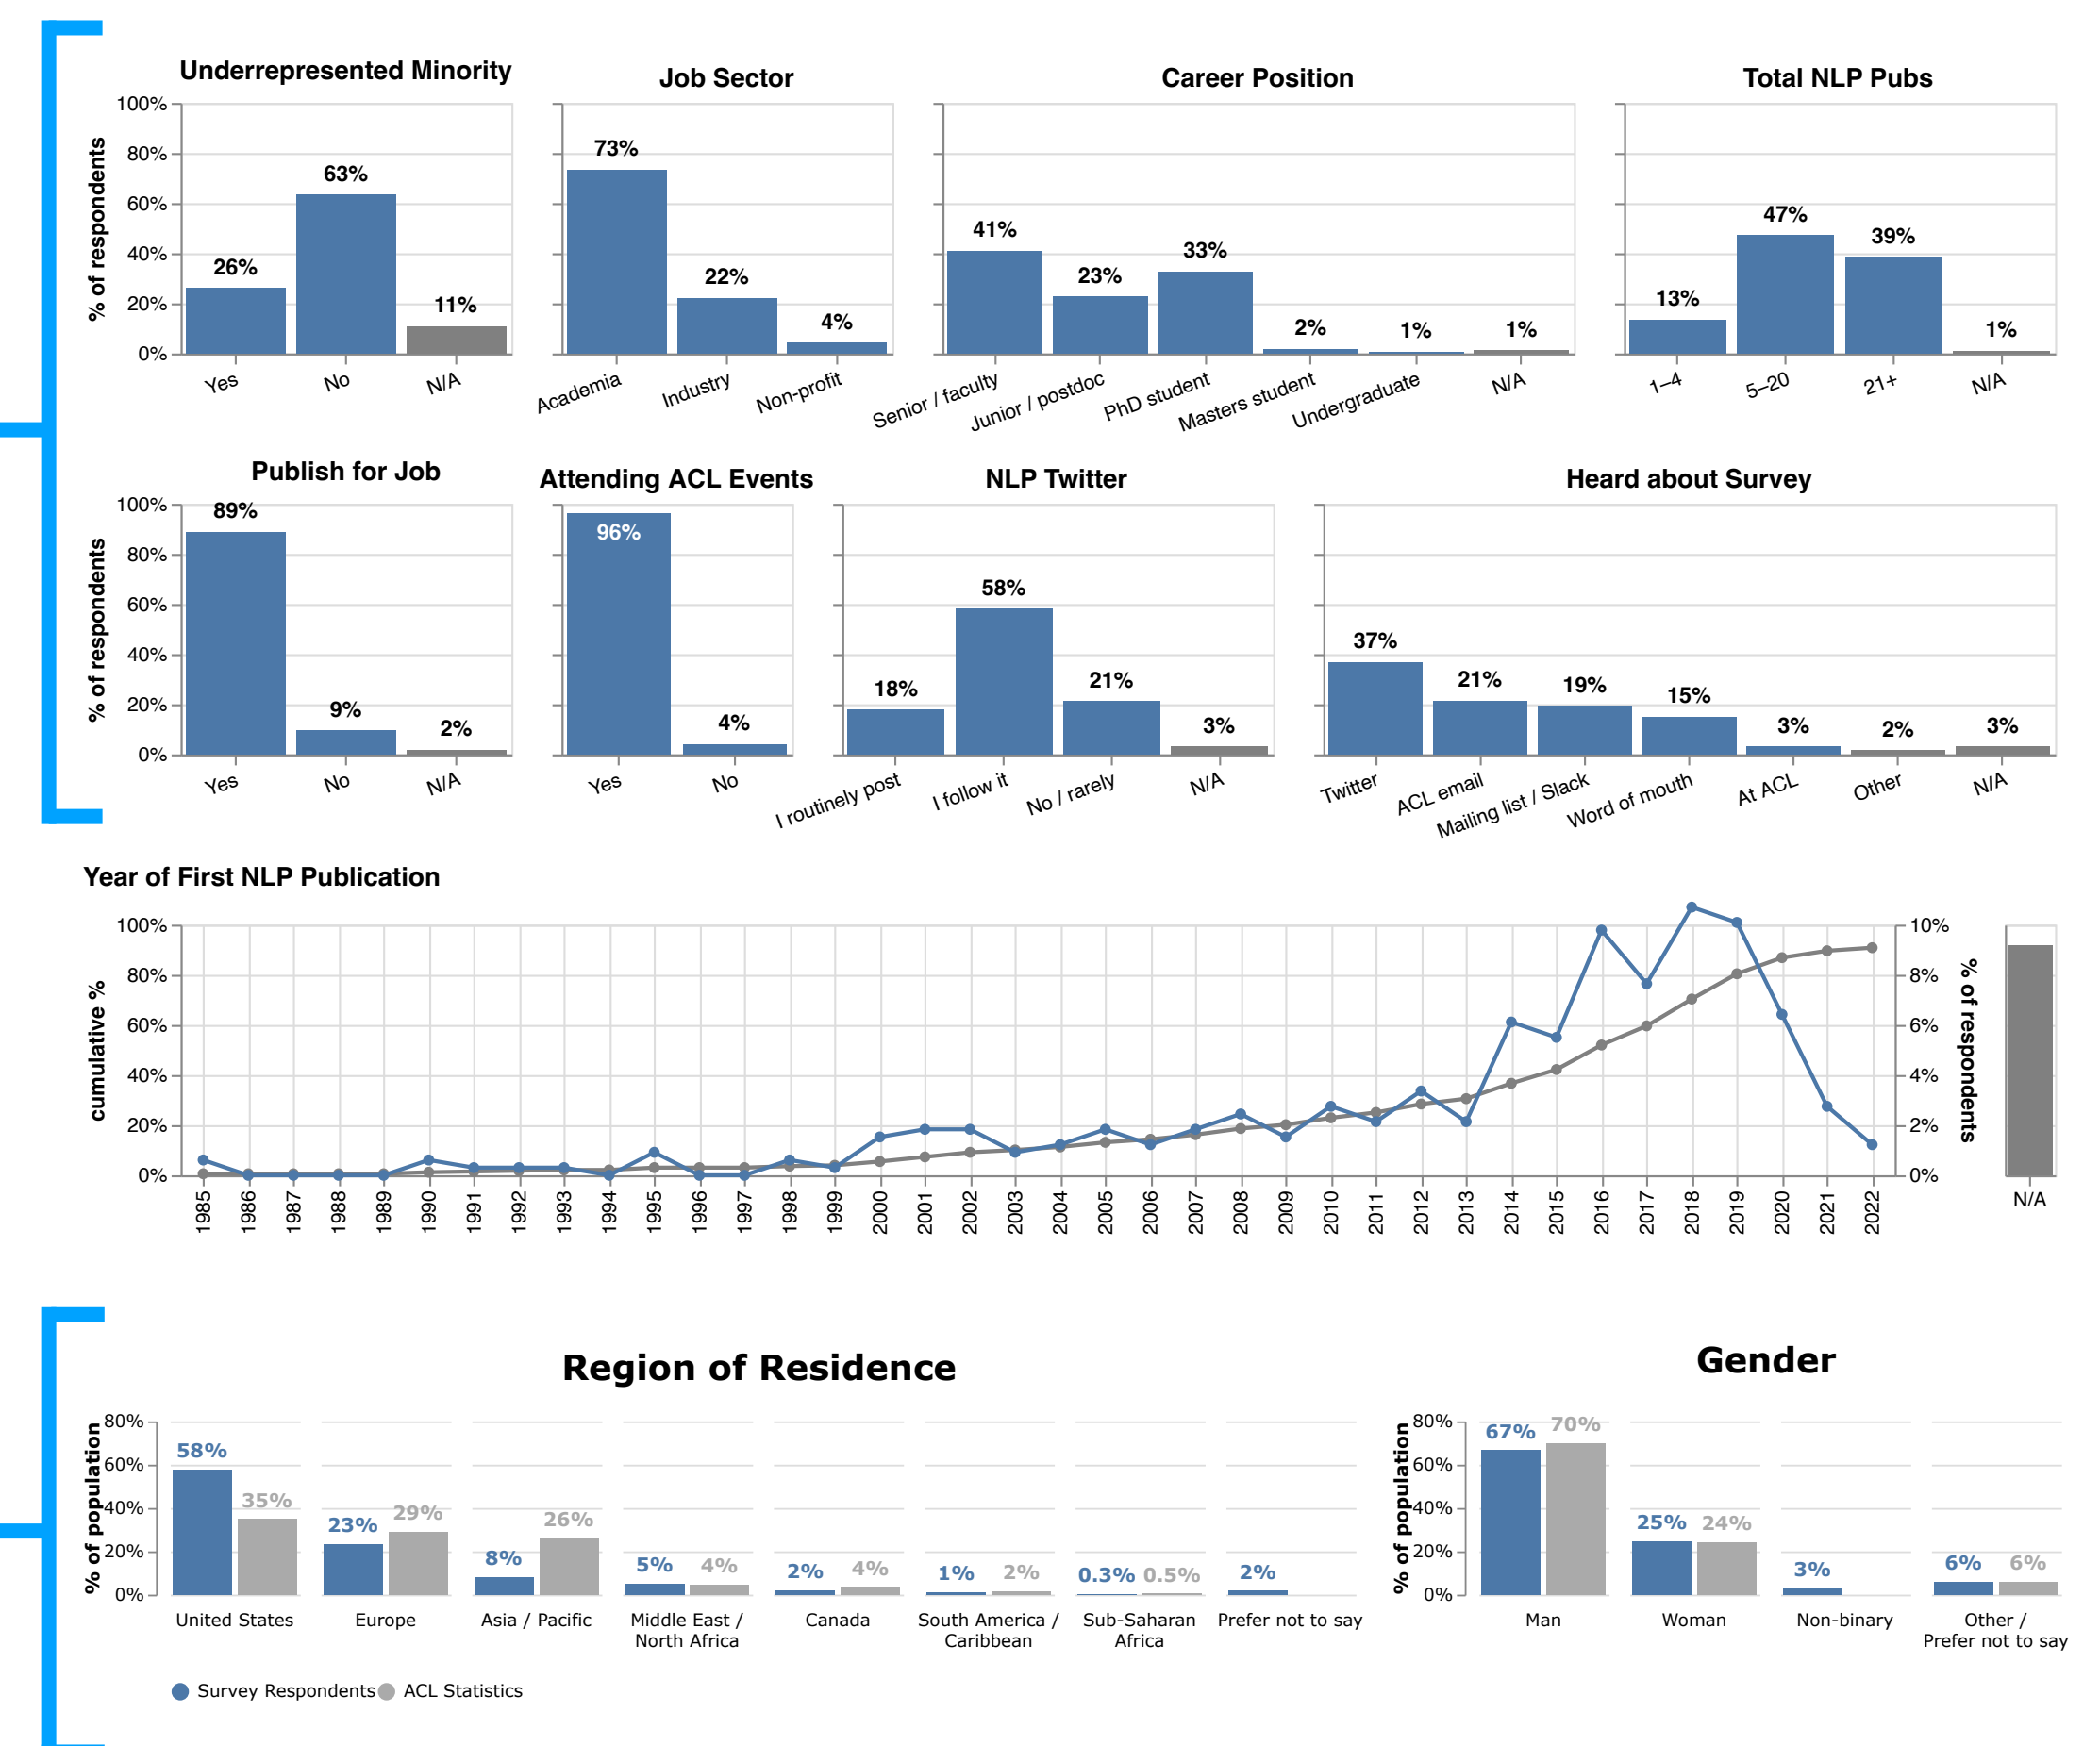
60% of respondents think carbon footprint is a major concern, but less than half think NLP should be regulated by governments.

Main Takeaways


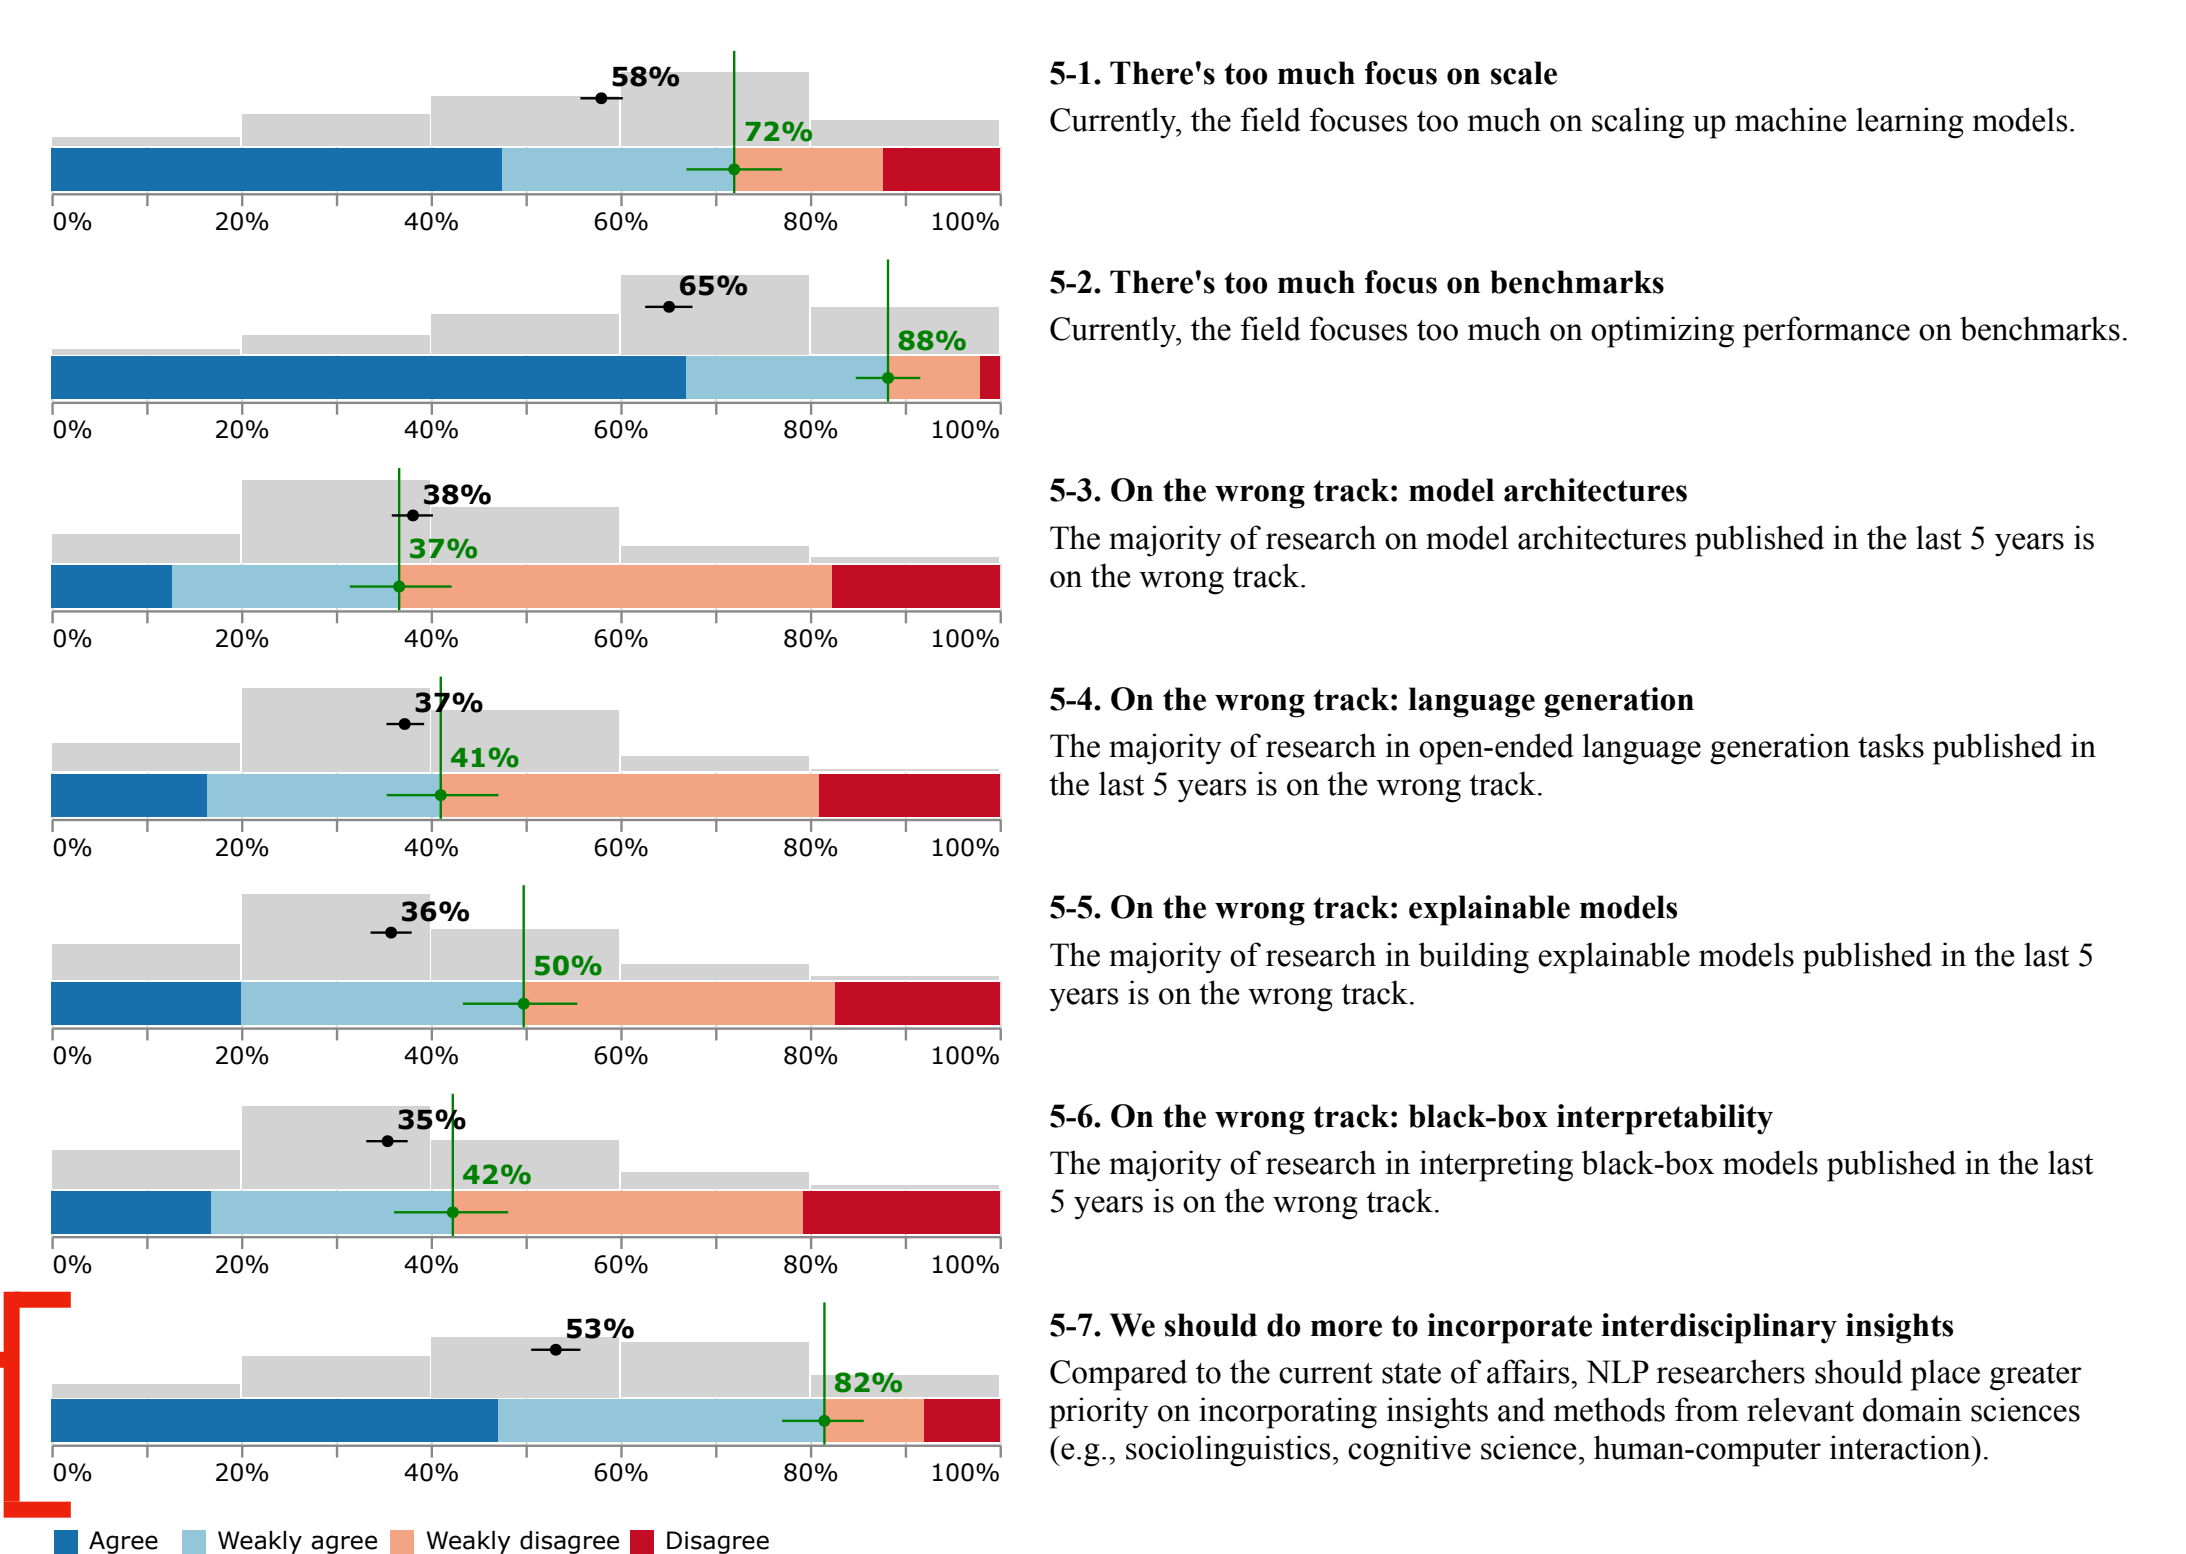
- Scaling maximalism was a surprisingly rare view in early/mid 2022
- Surprisingly many see a role for increased use of expert-designed inductive biases, linguistic theory, and interdisciplinary work
- There is little consensus on important ethical issues (e.g., misuse, environmental impact, regulation)
- A significant fraction (36%) see catastrophic risk from AI as plausible, and a majority (74%) see AI as plausibly transformative for society

Demographics

- 327 respondents in our target demographic (≥ 2 *CL pubs 2019–2022)



5. Promising Research Directions



6. Ethics

