

Pathway to Impact

6.S891/12.S992/6.S893: AI for Climate Action

Spring 2026

Speaker: Priya Donti

(incl. some content originally created with David Rolnick and Nsutezo Fobi for the Climate Change AI Summer School)

What does “impact” mean in AI for climate action?

Making direct progress on climate action: Reducing greenhouse gas emissions, improving adaptive capacity, improving understanding of climatic changes, etc.

Enabling some stakeholder to make better progress on climate action: E.g., improving decision-making, enabling new types of action, speeding up “time to action.”

Improving technical performance: Improving the performance of a model or modeling pipeline on a technical metric of importance to a climate-relevant problem.

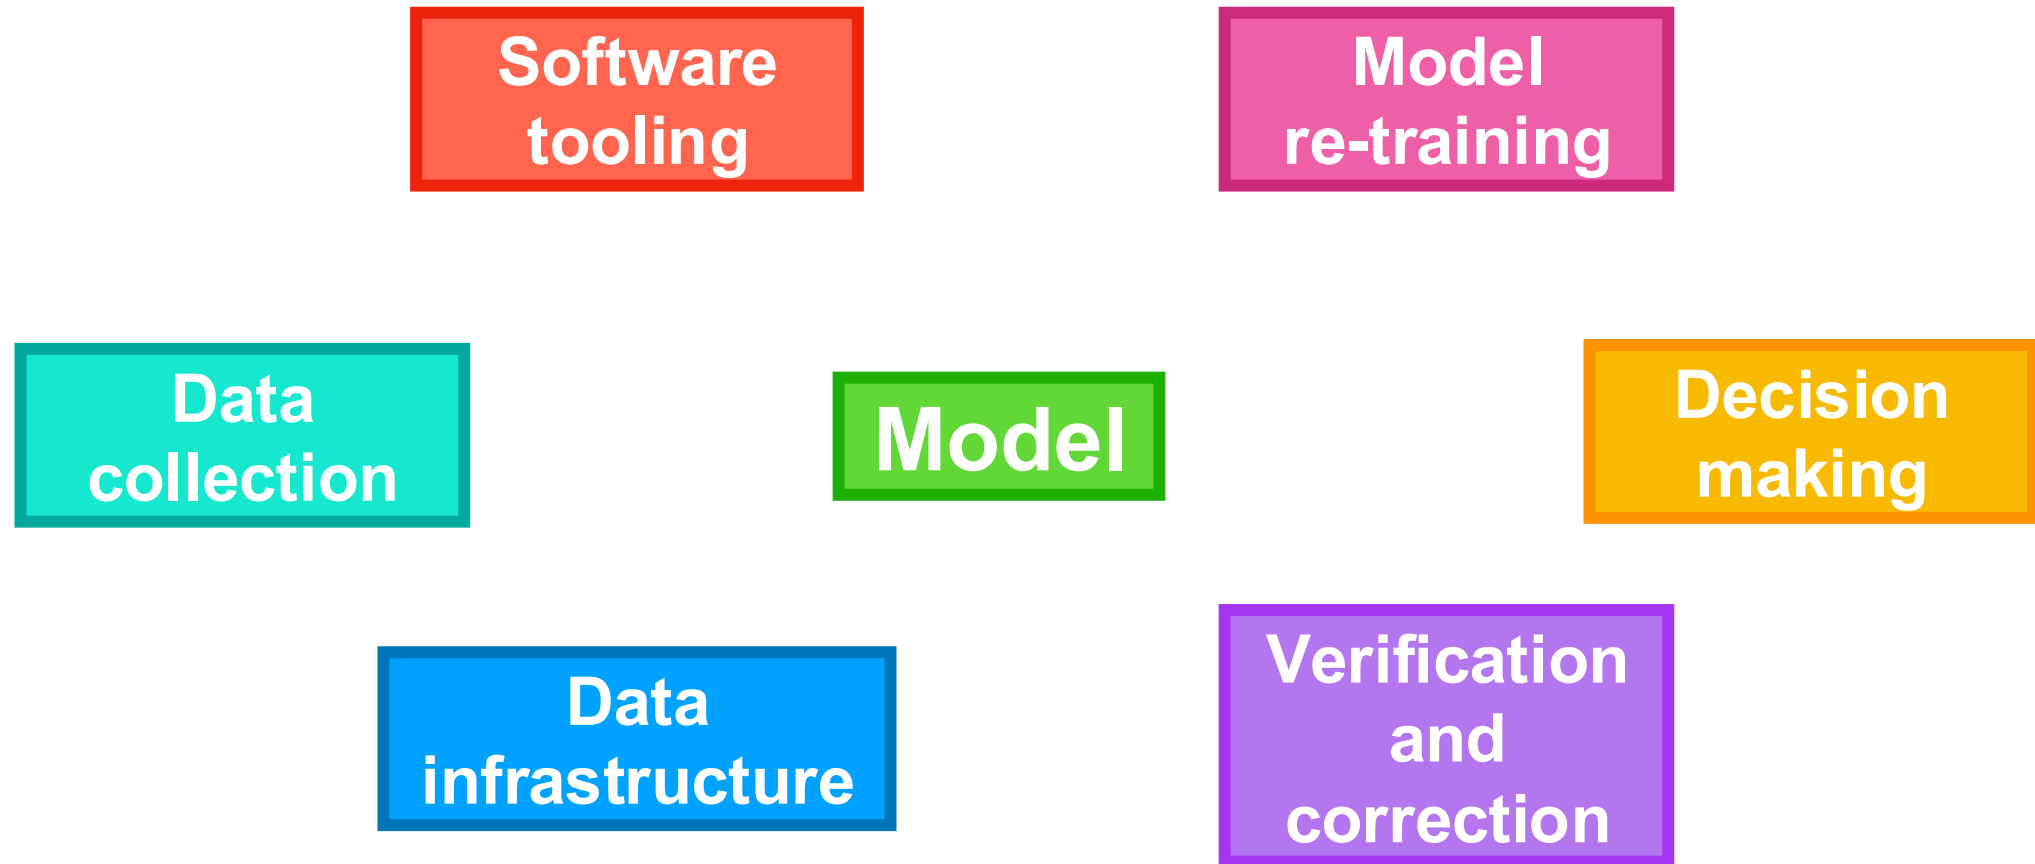
Improving the ability of others to contribute, and/or the quality of their contributions: E.g., creation of climate-relevant data and benchmarks, surfacing important problems and considerations, helping stakeholders “upskill.”

Avoiding adverse consequences: Not creating issues with respect to, e.g., negative environmental impacts, bias/fairness, access/participation, privacy/security, etc.

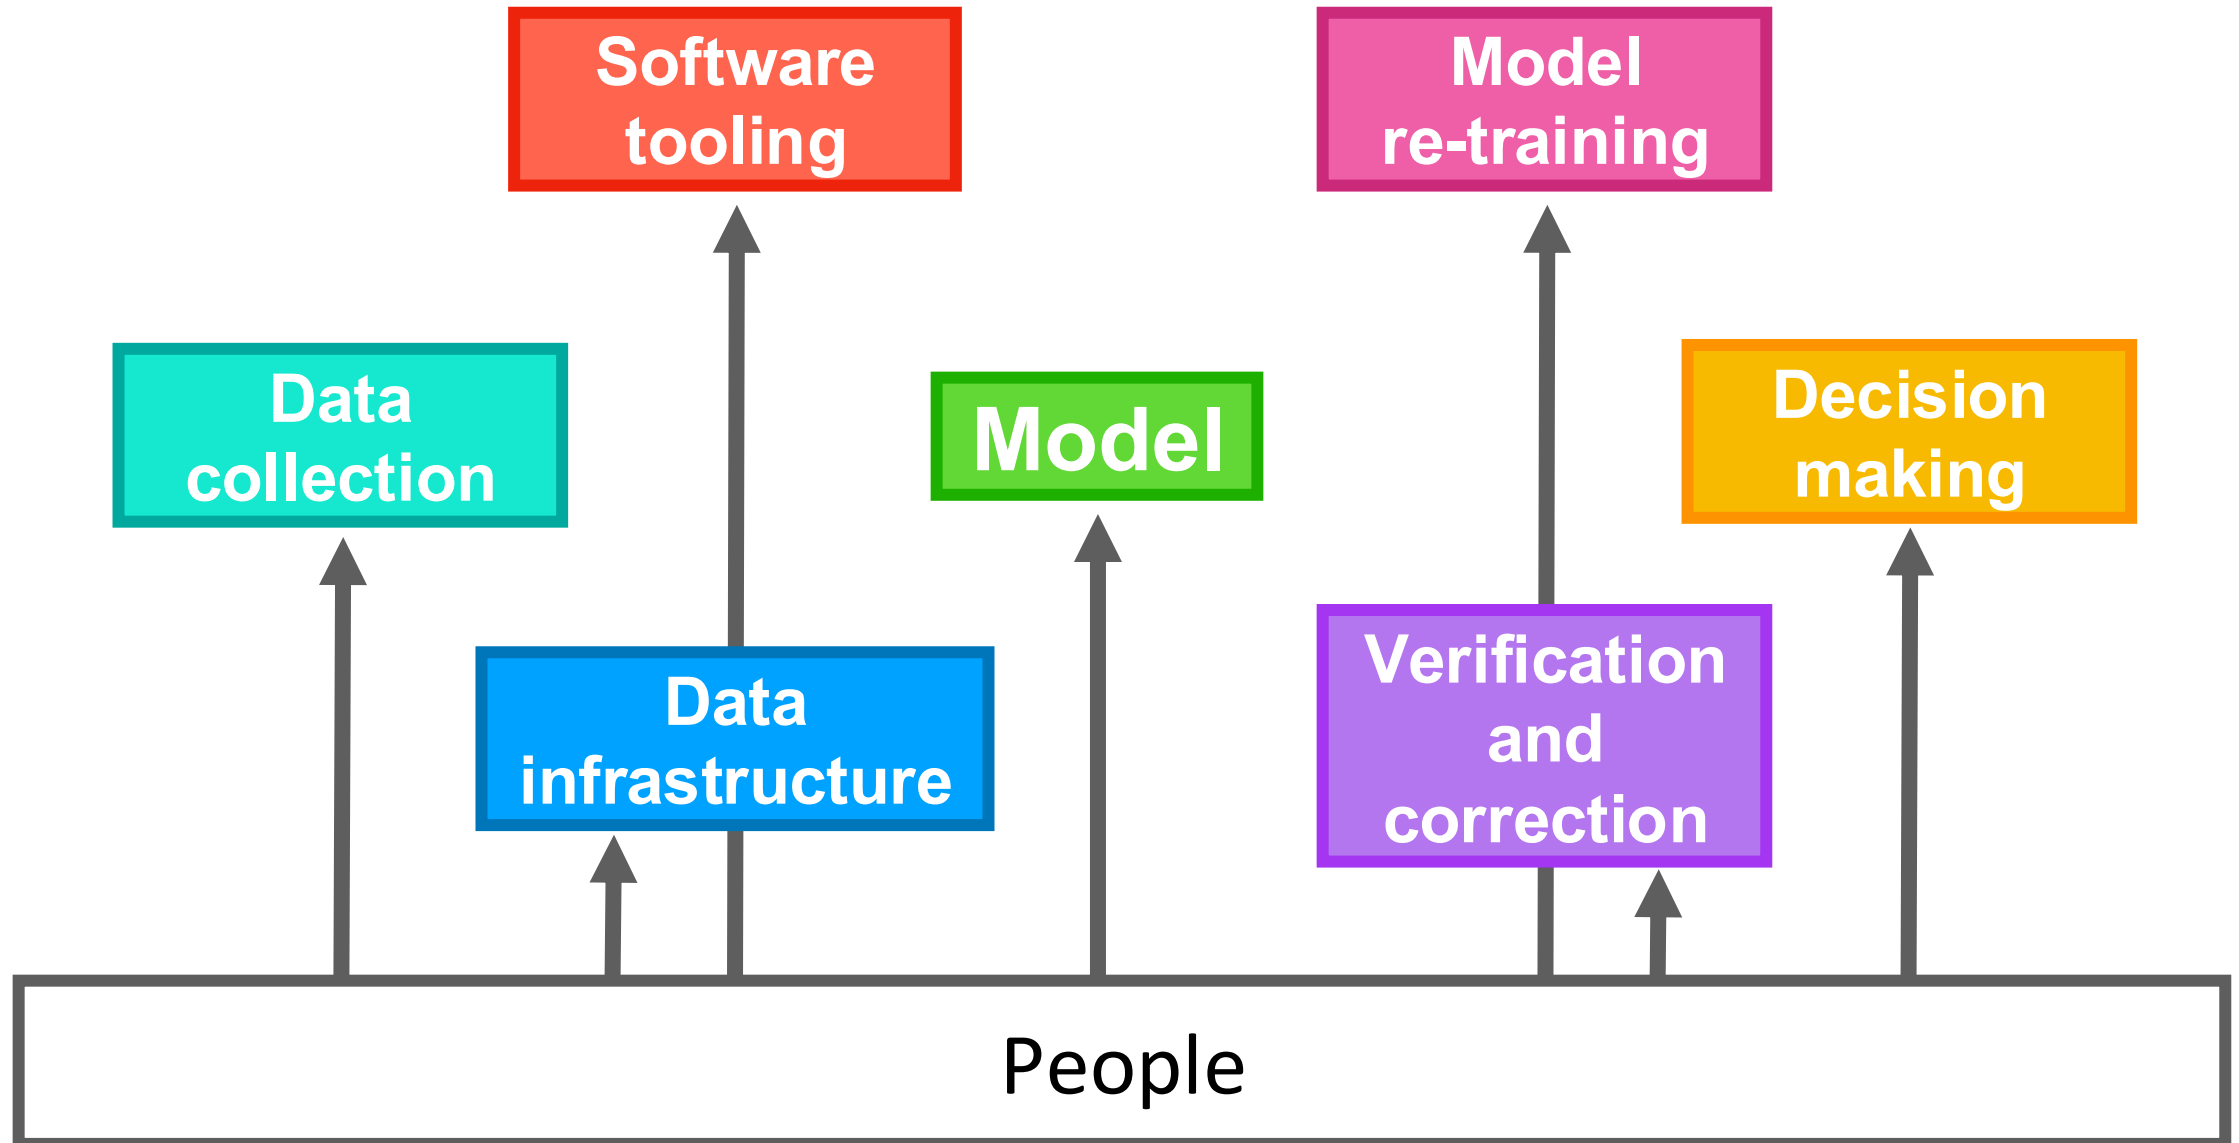
Recall: Good performance on realistic benchmarks \neq impact

Model

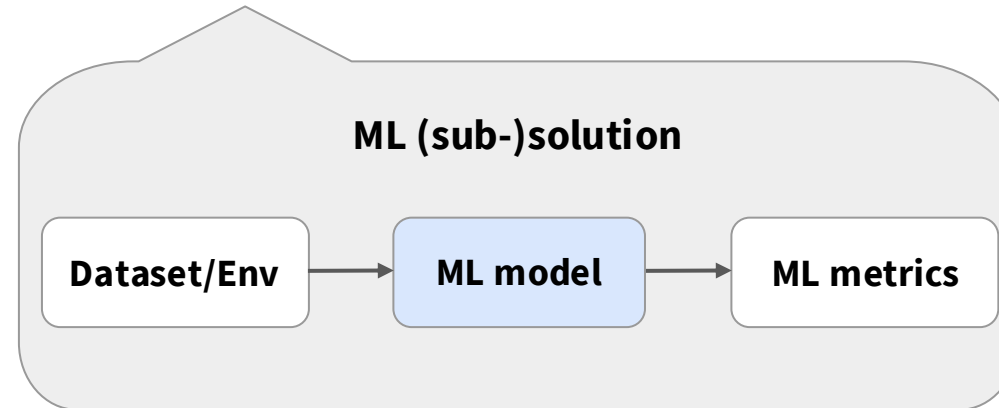
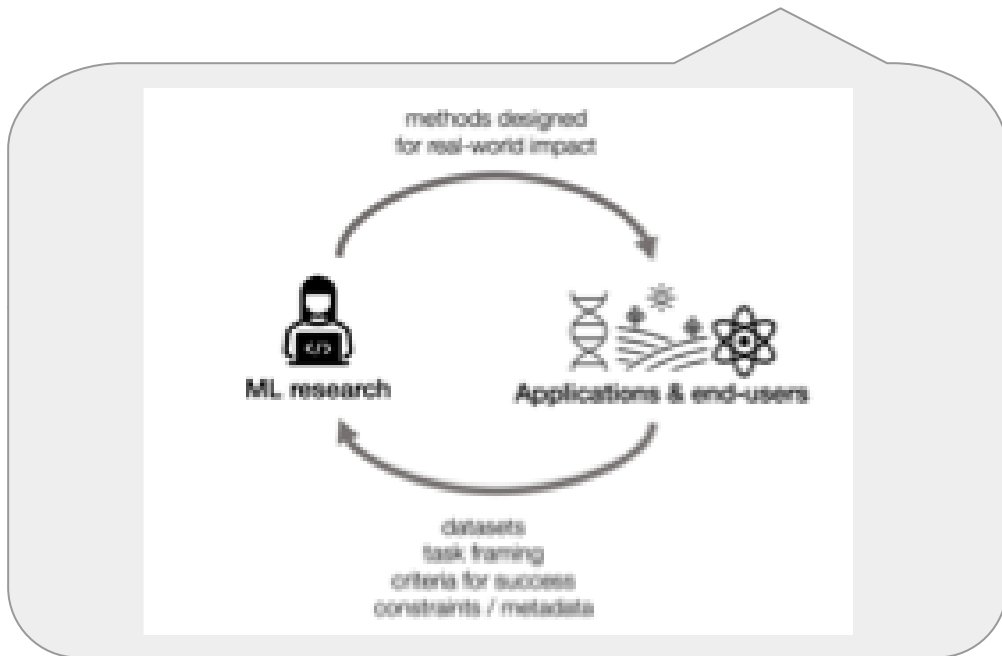
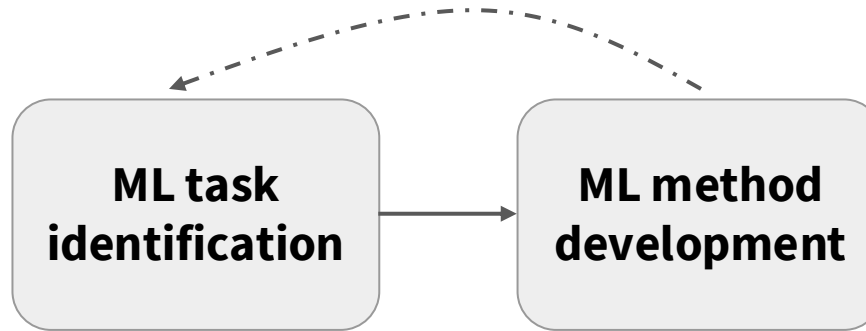
Recall: Good performance on realistic benchmarks \neq impact



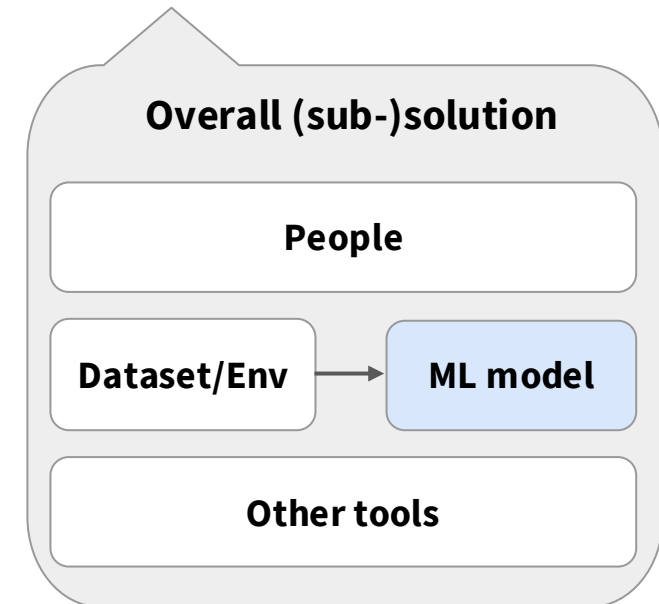
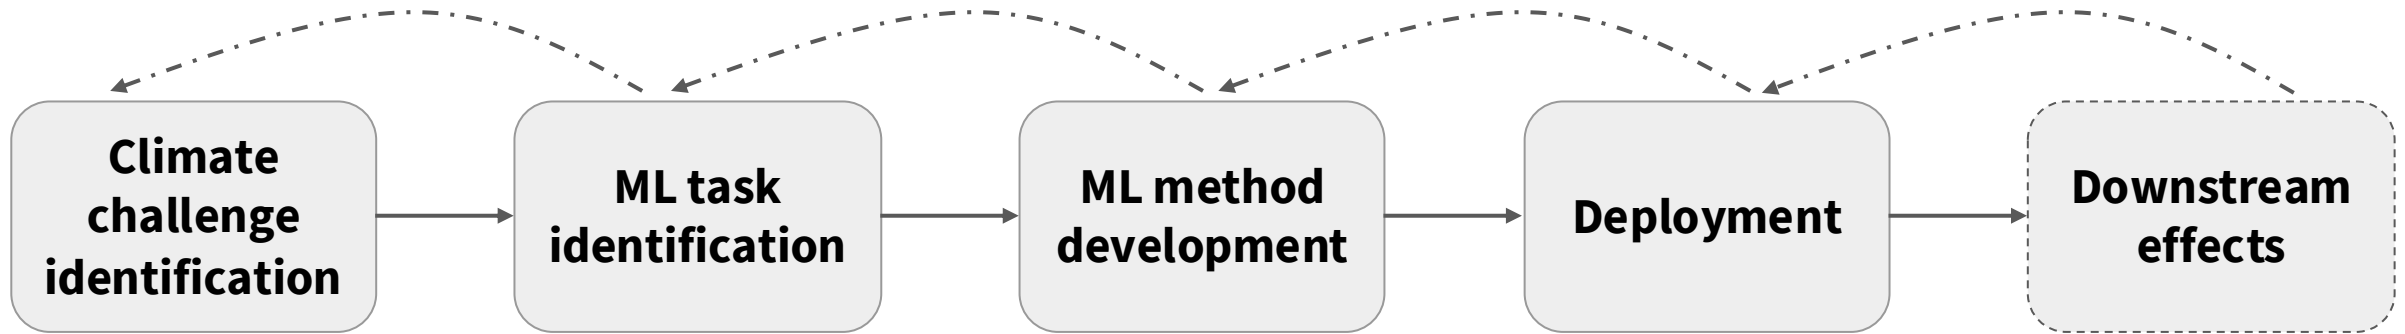
Recall: Good performance on realistic benchmarks \neq impact



Lifecycle of an ML-for-climate project

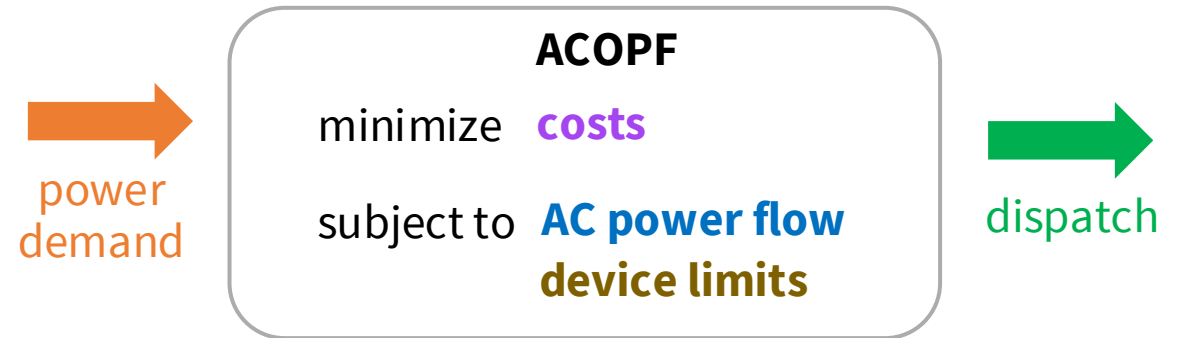


Lifecycle of an ML-for-climate project

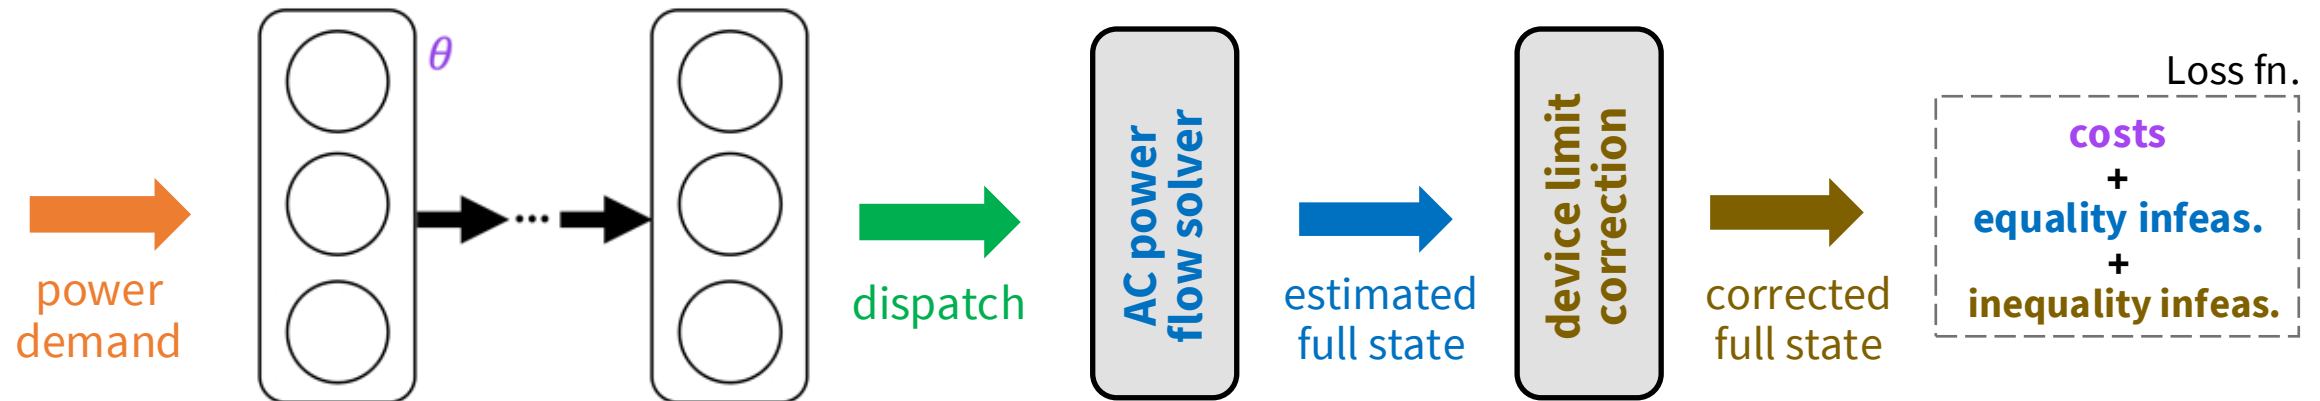


Running example: Surrogate models for grid optimization

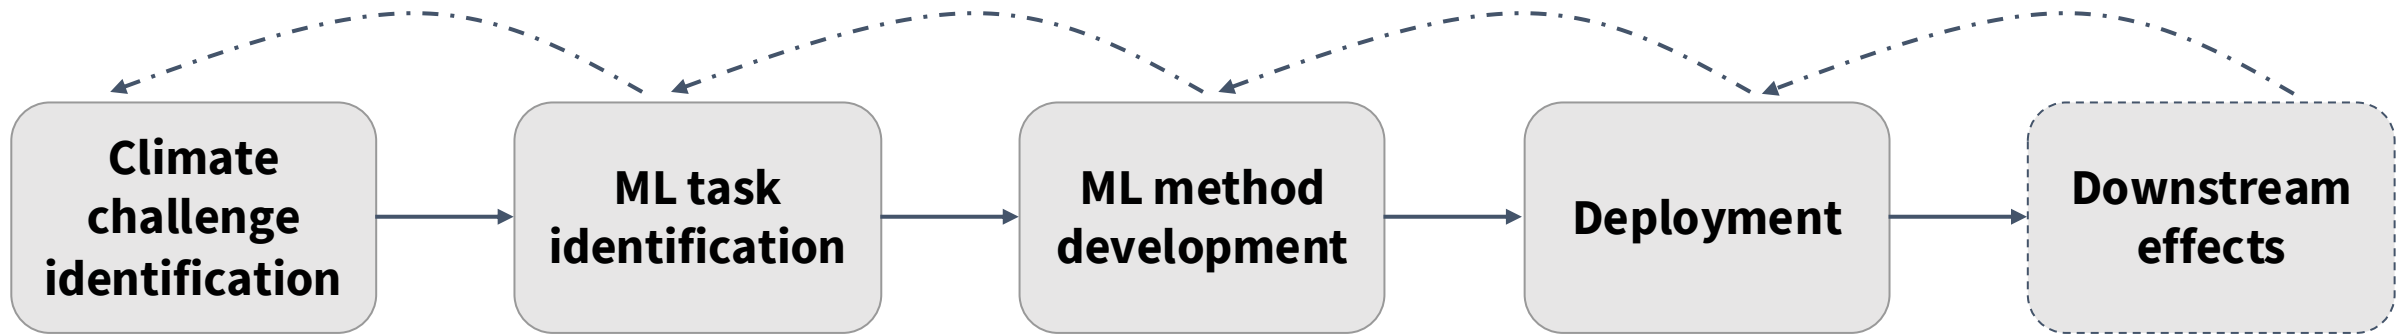
Goal: Provide fast, feasible approx. to AC optimal power flow (ACOPF)



Approach:



Example: Problem scoping



Today's processes for managing power grids are not sufficiently fast or scalable to enable large-scale integration of renewable energy

Speed up solution times for the ACOPF optimization problem

ML-based “surrogate models” that are fast, feasibility-preserving, and near-optimal

Deliverable: Algorithm (open-source)

- Inputs: Demand forecasts
- Usage: Standalone (human user) or integrated with other grid optimization algorithms, for operations and markets

Relationship: ACOPF is a foundational problem for grid management, used alongside other grid analysis/prediction/optimization/control tools and within power markets, via both automated and human-in-the-loop processes

Simplifications: We assume fixed power grid topologies, perfect supply/demand forecasts, and perfect knowledge of prices

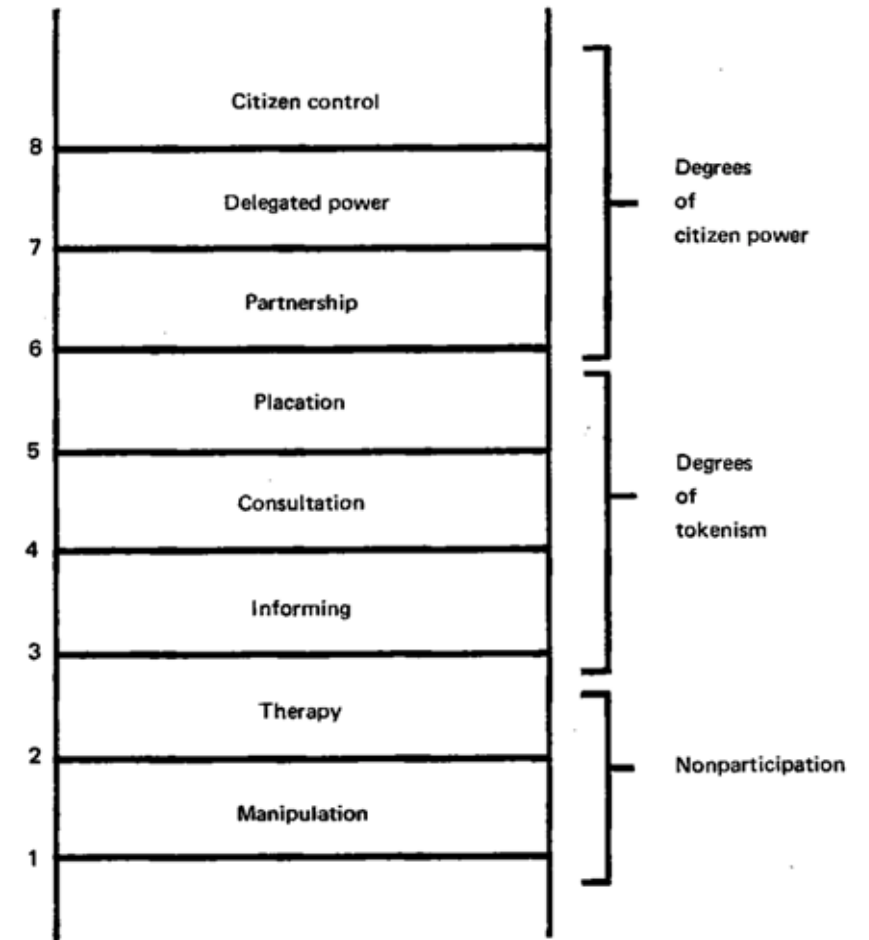
Mapping and engaging project stakeholders

Stakeholder types (e.g.)

- Researchers (tech & social sciences)
- Implementing entities and industries
- End users
- Policymakers
- Other affected parties

Going beyond the “hypothetical stakeholder”

Meaningful engagement required



Arnstein's Ladder of Citizen Participation

Guide to collaboration and partnership

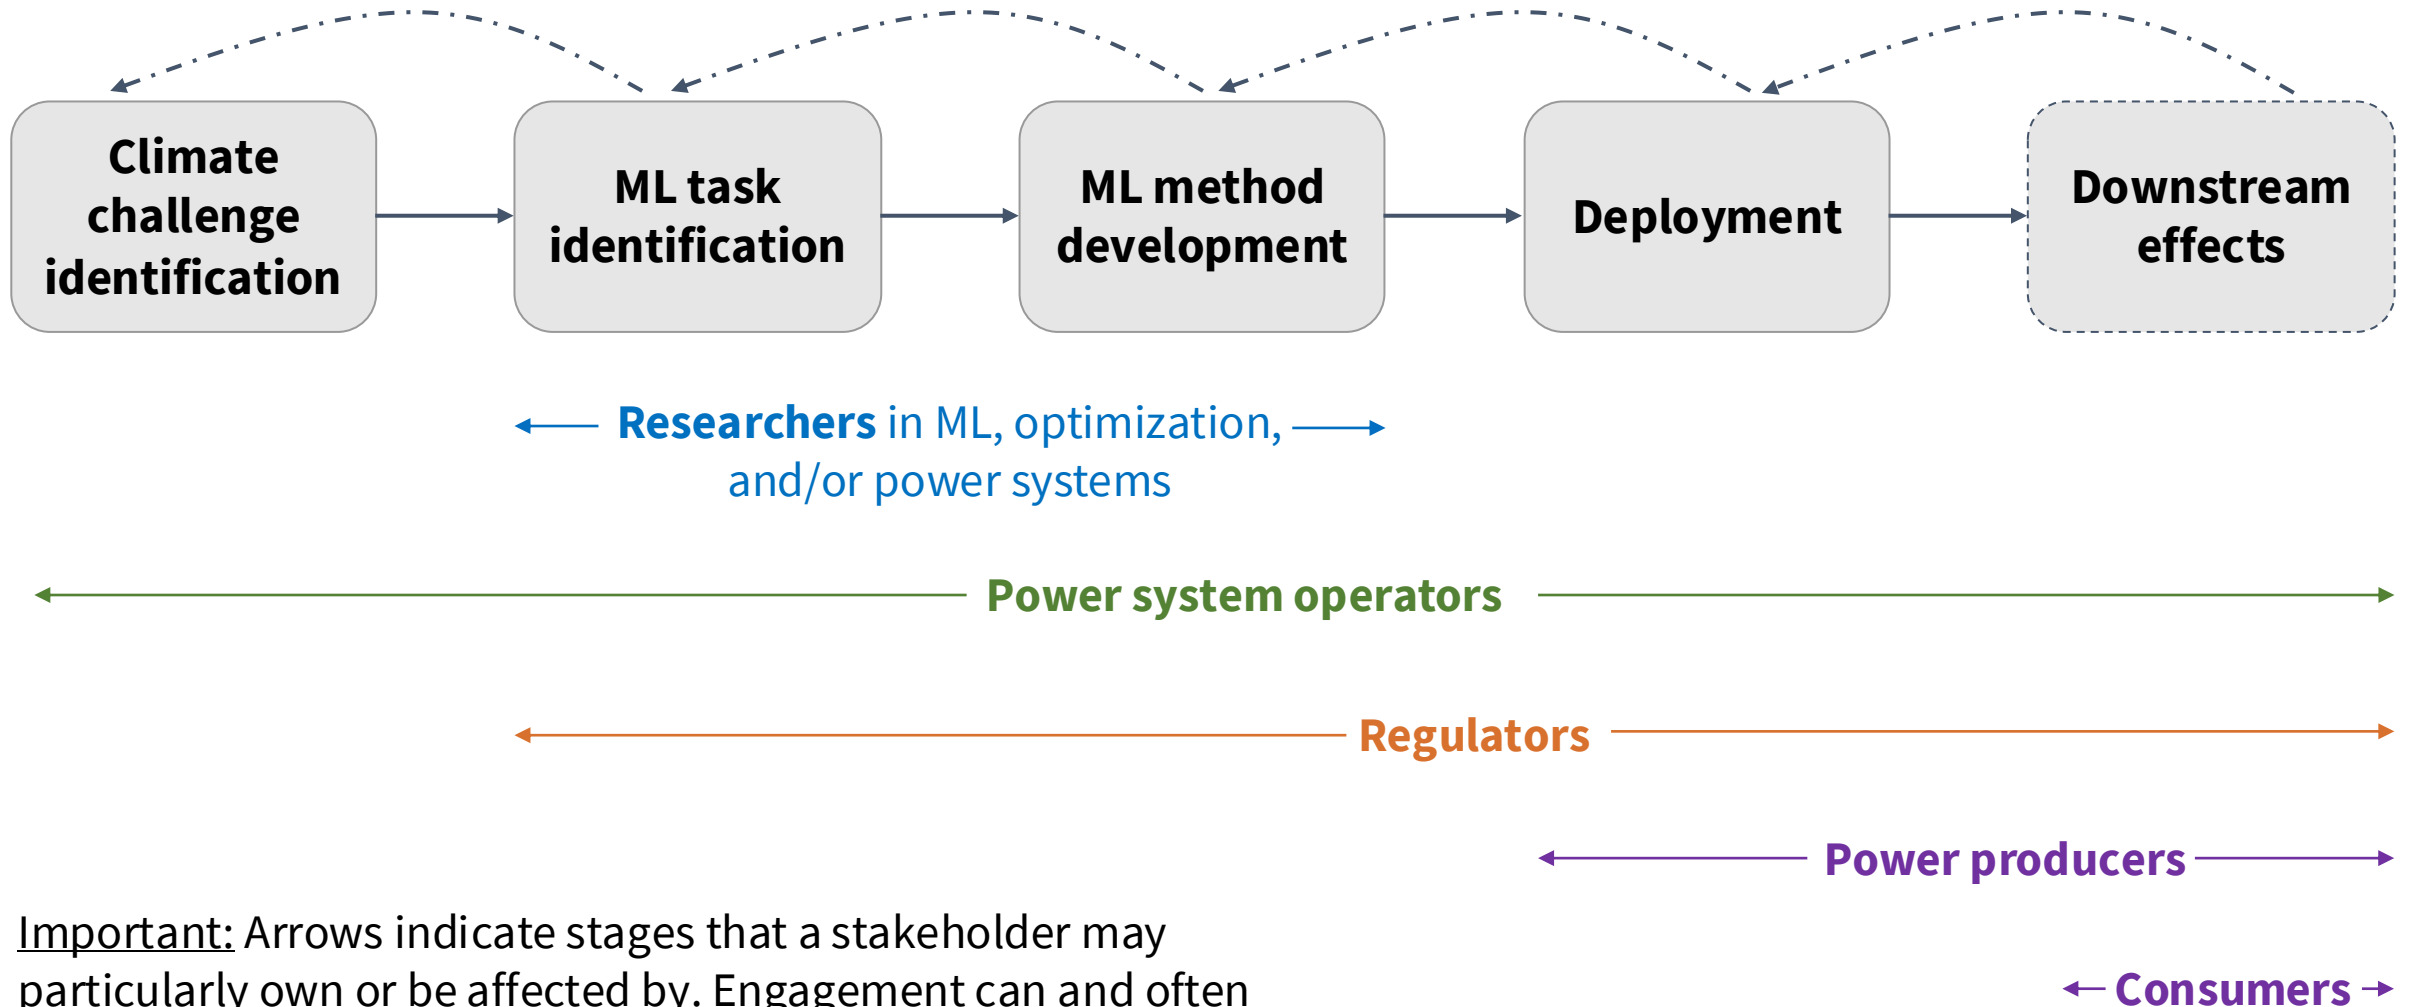
How to find collaborators:

- Consult the literature - academic, technical reports, policy docs, etc.
- Attend workshops, conferences, seminars, etc. - talk with speakers & participants
- Look at, e.g., industry consortia that may have partnership initiatives
- Consider if your own org has an office or support for tech transfer
- Research relevant groups at your own org, and discuss with members of your own org who may have relevant partnerships

Involve collaborators as early as possible, to help in

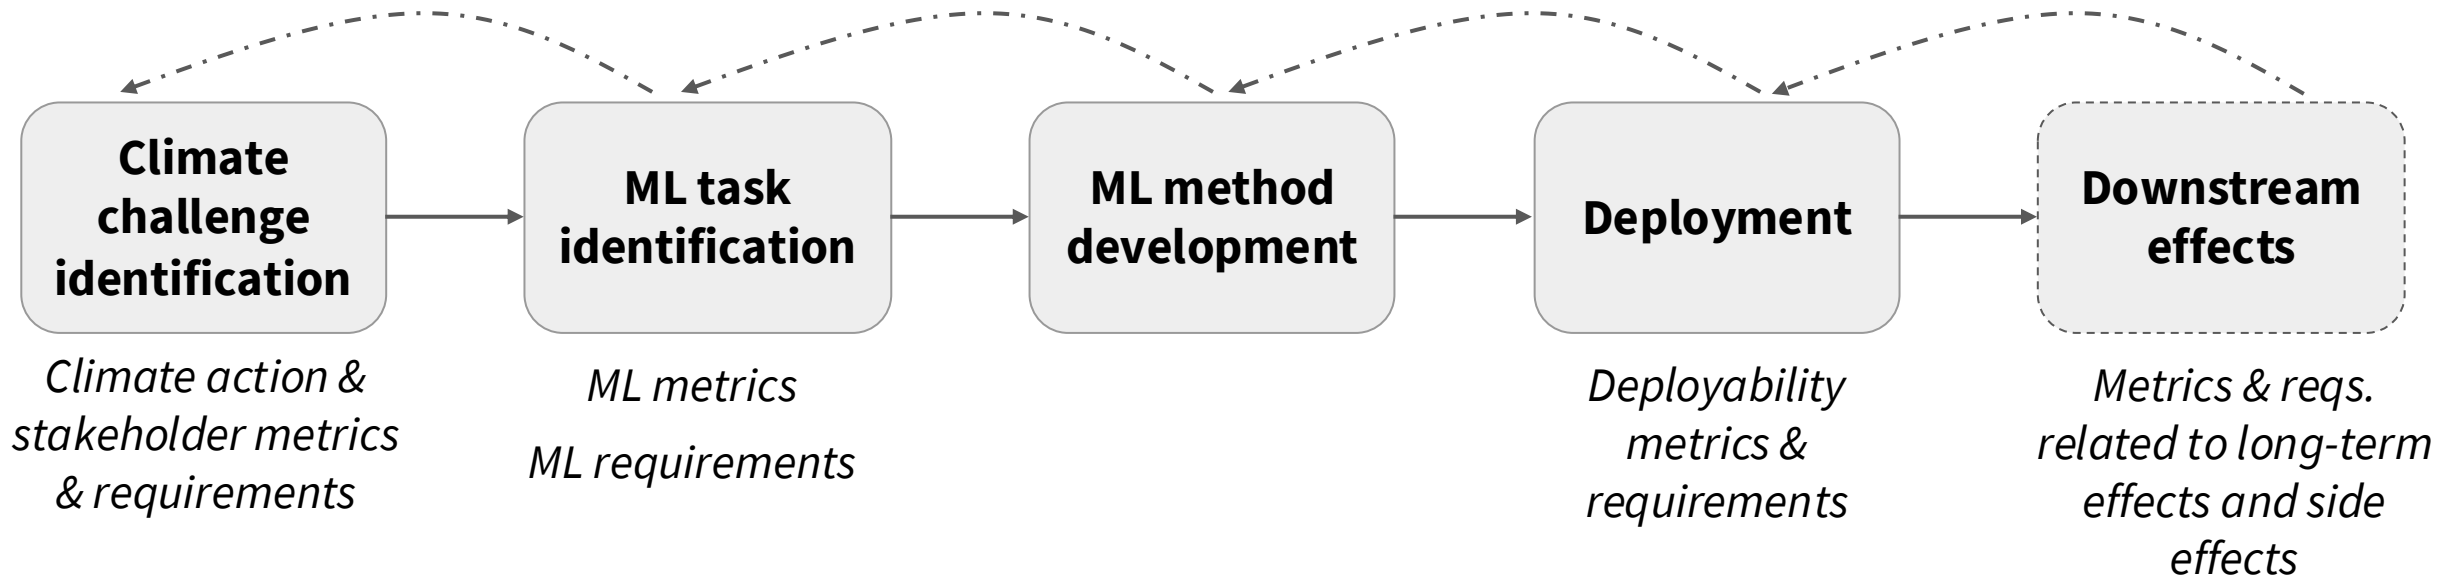
- Problem framing
- Identification of important constraints, pitfalls, or metrics for success
- Incorporating domain knowledge
- Shaping deployment

Example: Stakeholders

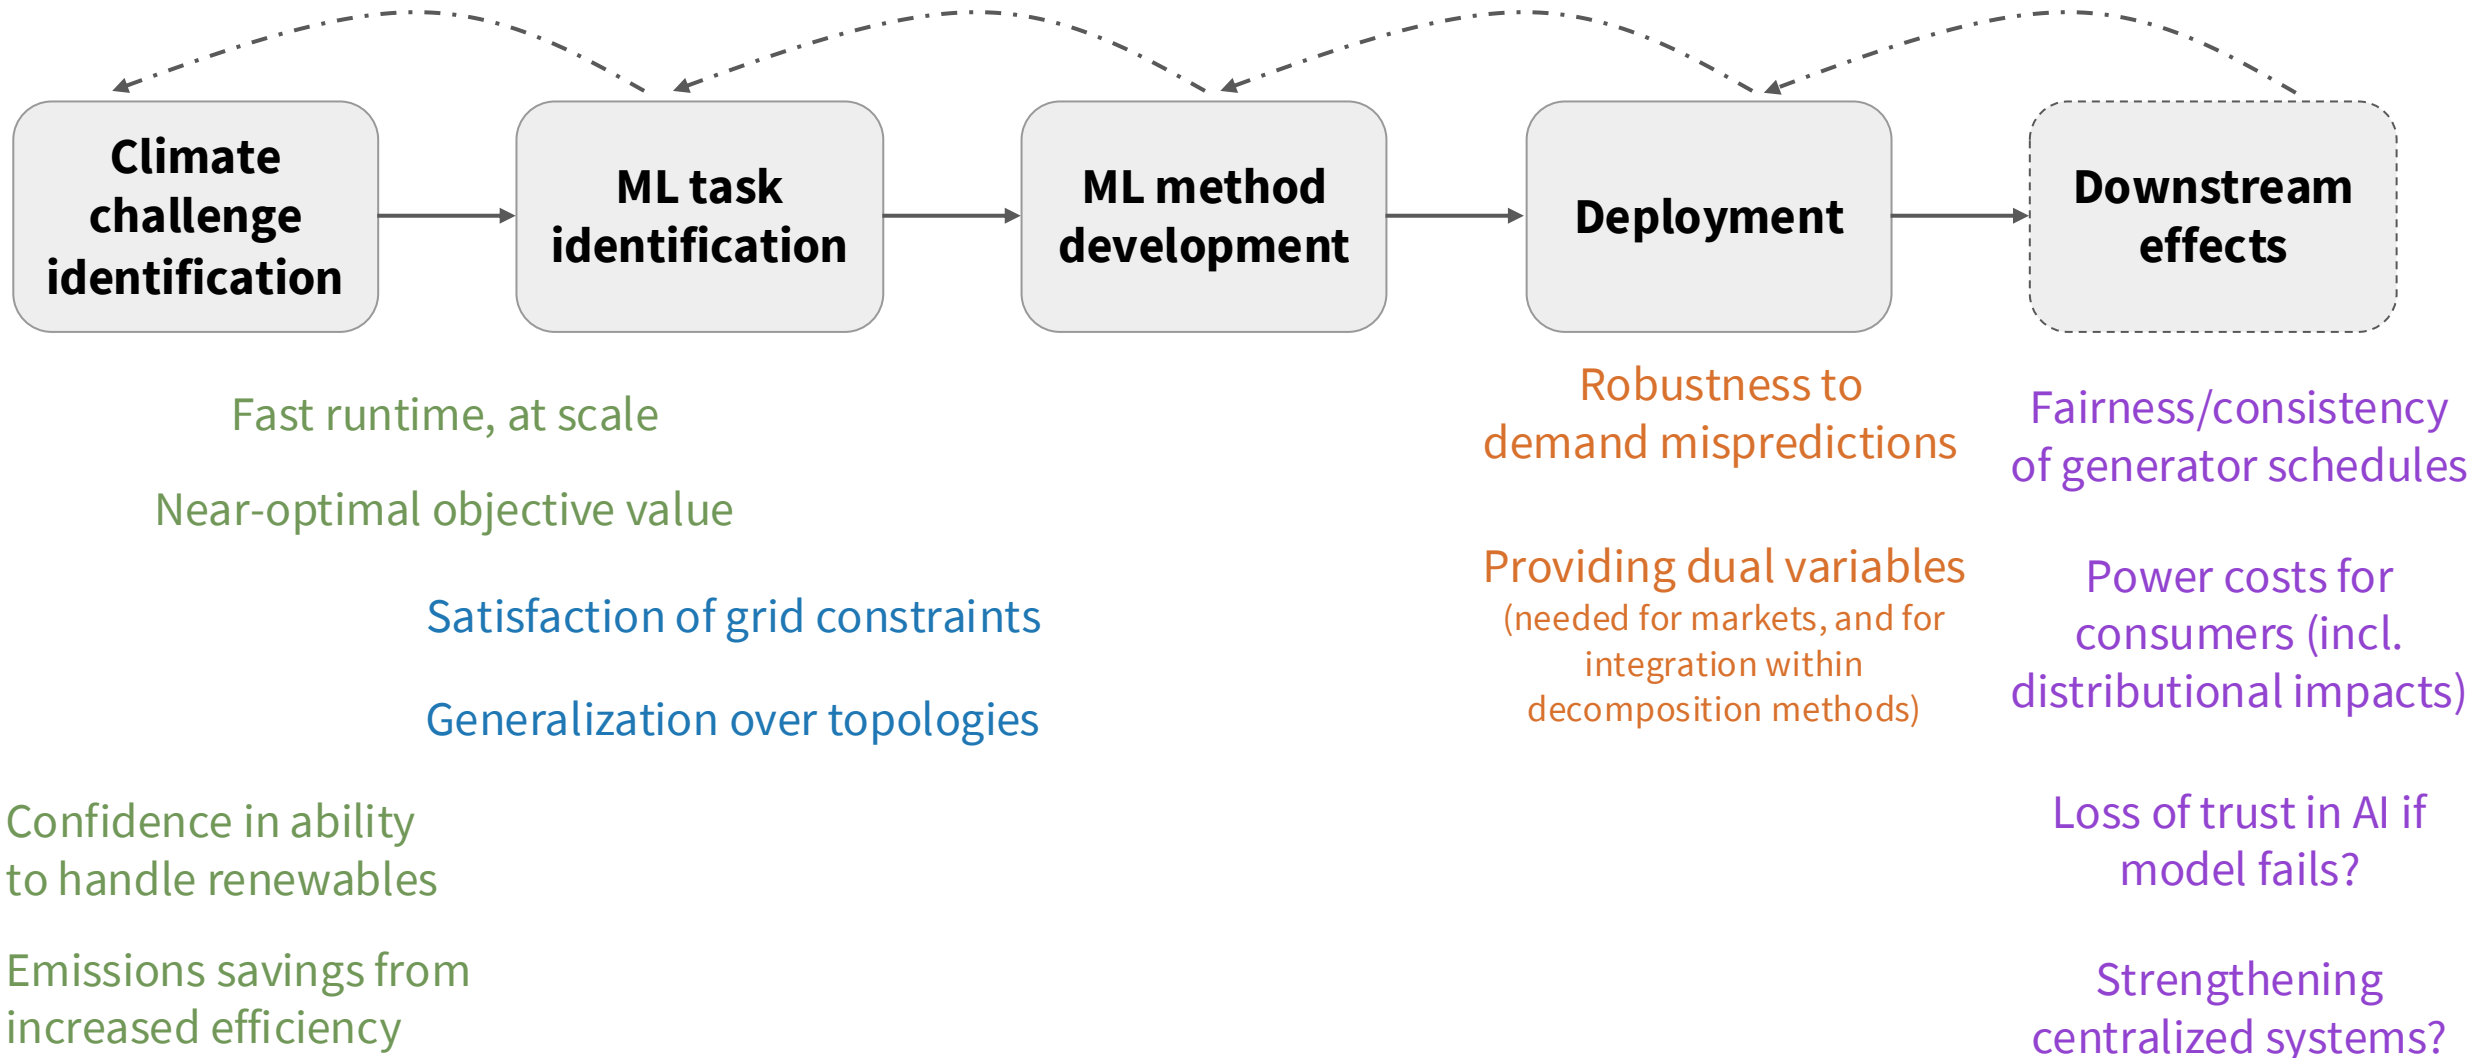


Important: Arrows indicate stages that a stakeholder may particularly own or be affected by. Engagement can and often should start earlier than that!

Requirements and evaluation



Example: Requirements and evaluation



Long-term effects

What might happen if the solution is deployed and continues to be deployed, which might not be captured by “instantaneous” evaluation?

Example considerations:

- Rebound effects: If something is made more efficient, will it be used more? Will this partially or completely counteract the gains from efficiency?
- Feedback loops: If the model is used, how does that change future outcomes? How does that, in turn, change model inputs? Does that, in turn, change the distribution of outcomes over time?
- Effects of process changes: How does the deployment of the algorithm change the way that things are done (e.g., how systems are planned or envisioned)? What are potential effects associated with that?

Responsible AI for climate action

Mitigating biases in data and models

- ▶ E.g., Buildings data: Housing discrimination, geographic disparities in availability
- ▶ E.g., Weather models: Calibration may be optimized for particular regions

Improving trustworthiness and accountability

- ▶ Safety and robustness: Critical in, e.g., power systems and industrial operations
- ▶ Interpretability and auditability: Critical in, e.g., policymaking contexts
- ▶ Privacy preservation: Critical, e.g., when working with personal data

Centering equity and climate justice

- ▶ Centering diverse stakeholders: E.g., industrial ag vs. smallholder farmers
- ▶ Avoiding centralization: Democratized capacity and compute, digital divide
- ▶ Avoiding digital colonialism: E.g., smart meters, analysis of remote sensing data

Sustainability impacts

- ▶ Model emissions: Operational and hardware impacts, across the project lifecycle
- ▶ Effects on other sustainable development goals (SDGs): Co-benefits, harms

Accounting for potential “dual use”

Technology Readiness Levels (TRLs)

Framework for assessing the maturity of a particular technology

Additional forms of readiness

- Market readiness level (MRL)
- Adoption readiness level (ARL)
- Societal readiness level (SRL)
- *Etc.*



Advancing readiness

Example bottleneck	Example ways forward
Quality of algorithmic output	Improving ML methods
Not satisfying on-the-ground requirements (e.g., interpretability, uncertainty quantification, various notions of robustness)	Improving benchmarks, data/simulators, and evaluation
Dataset quality, size, and availability, or limited geographic scope	Enhancing data collection Data augmentation
Lack of computational infrastructure to support deployment (e.g., simulators, ML hardware)	Improving access to computational infra. Computationally efficient ML methods
Lack of knowledge or capacity in relevant deployment orgs	Training and upskilling Identifying additional staff/capacity Improving usability and user interfaces
Trust gaps in relevant stakeholder communities	Meaningful stakeholder engagement
Lack of regulation to provide incentives or legal structure surrounding deployment	Engagement with regulators and policymakers Clever business models
Negative side effects	Modifying how the model is used/interpreted Improving measurement and transparency

