



Deploying AI4SG Applications

From Optimization to Sustainable Change

Amulya Yadav

Associate Professor, College of IST, Penn State

A decade of technically sophisticated interventions



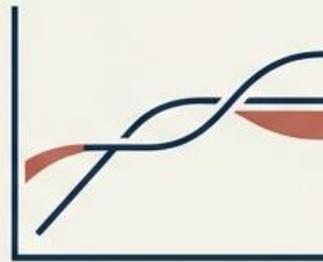
HIV Prevention AI

Predictive models and adaptive policies for targeting and resource allocation.



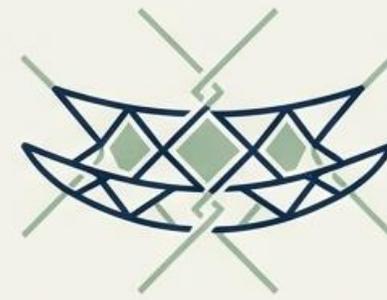
Substance Abuse AI

Risk prediction and continuous engagement data optimization for intervention efficiency



Suicide Prevention AI

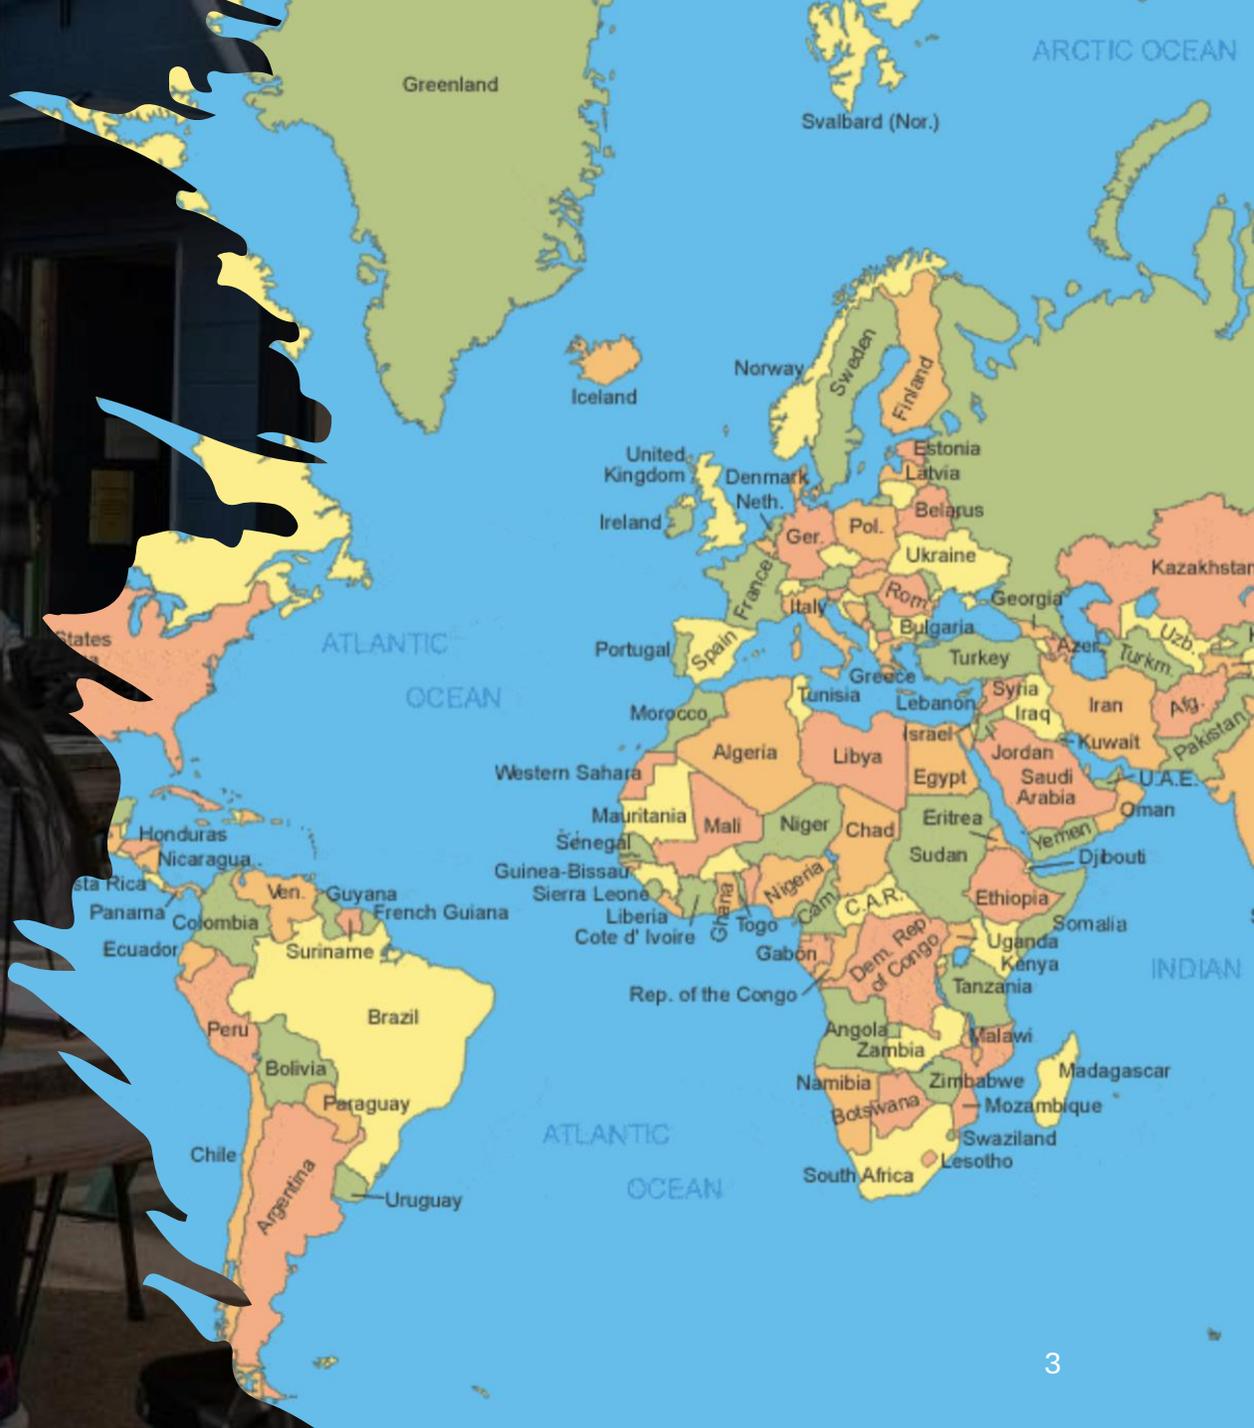
Deep risk targeting layered onto existing clinical service frameworks.



All built with real partners, addressing real needs, and featuring strong technical validation.

Case Study 1

AI for HIV Prevention



Case Study 1: Description

- NGO Collaborator:



- *What do they do?* Run Peer Led Popular Opinion Leader (POL) interventions to spread awareness about STDs and to change behaviors
- *What did they need help with?* How to choose peer leaders for this intervention
- *How did we help?* HEALER and CHANGE: Two AI agents that could identify key peer leaders who would serve as good “influencers” in the network (*JAIDS, AAAI 21, IJCAI 18, AAMAS 18a, AAMAS 18b, AAMAS 17a, AAMAS 17b, IJCAI 17, AI Magazine, AAMAS 16, IAAI 15*)

Case Study 1: High-Level Explanation



Case Study 1: Key Results

Field deployment

Trial with 718 homeless youth over 2 years

Collaboration with three Los Angeles drop-in centers; USC School of Social Work

Significant improvement over status quo

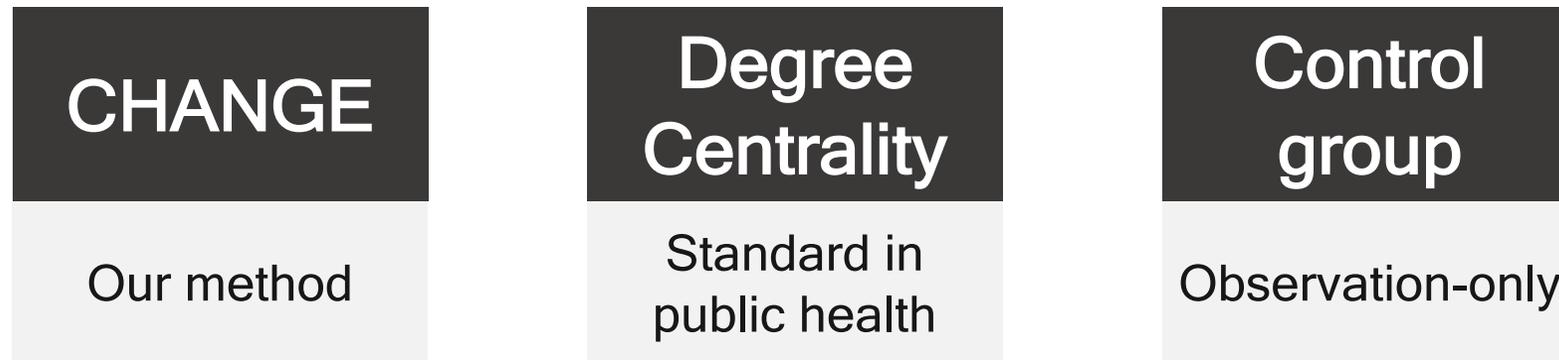


Conducting surveys at SPY



Case Study 1: Key Results

Trial design: three arms



Case Study 1: Key Results

Each deployment



Baseline

Recruit 80 youth
Survey baseline behaviors



Intervention

Train 10-12 peer leaders



Follow-up

Survey all 80 at 1 and 3 months
Behavior change?

Case Study 1: Key Results

Each deployment



Baseline

Recruit 80 youth
Survey baseline behaviors



Intervention

Train 10-12 peer leaders

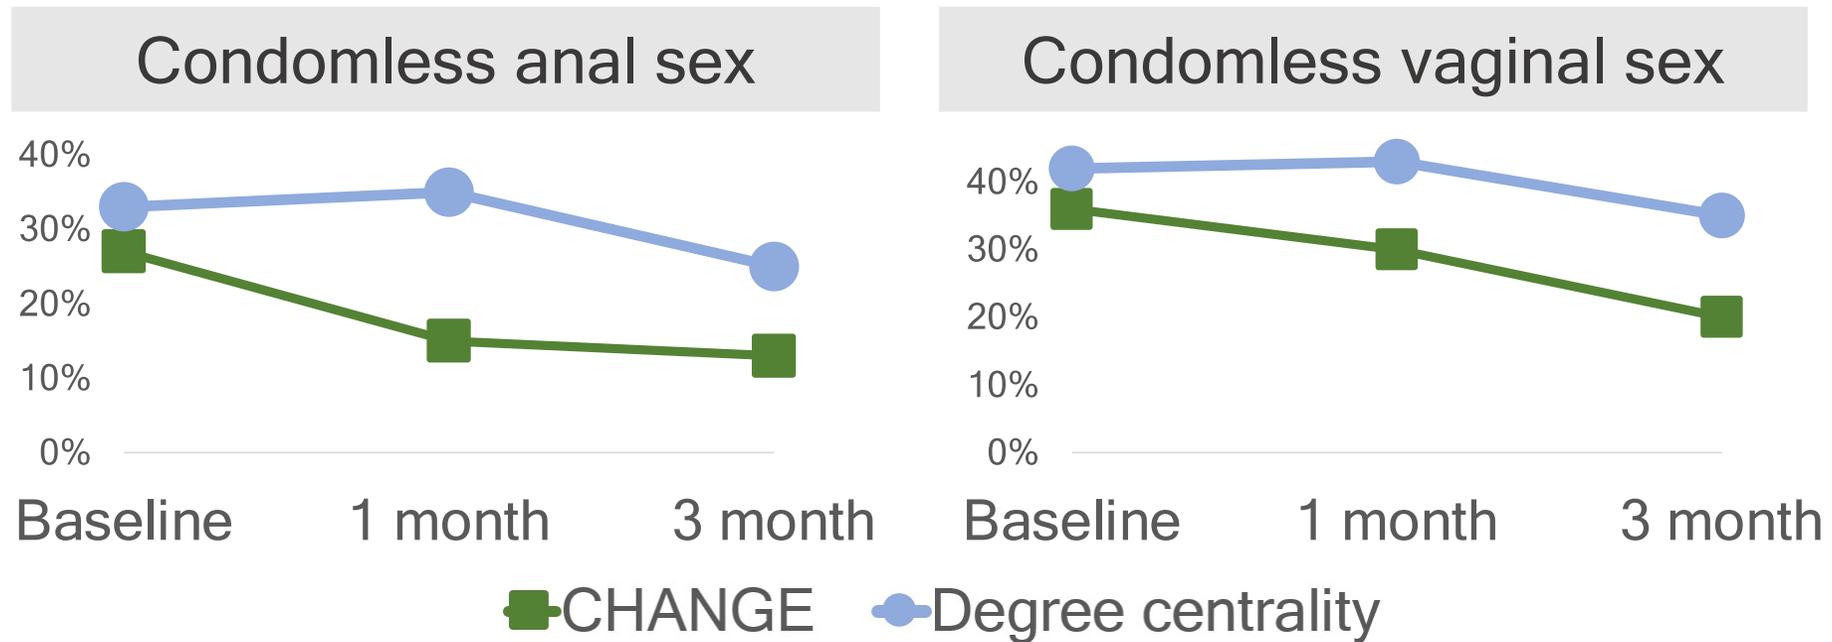


Follow-up

Survey all 80 at 1 and 3 months
Behavior change?

Case Study 1: Key Results

Results



Wilder et al. Clinical Trial of an AI-augmented Intervention for HIV Prevention in Youth Experiencing Homelessness. AAAI 2021.

Faster and larger decrease in CHANGE group

Case Study 2

AI for Substance Abuse



Case Study 2: Description



UNIVERSITY of
DENVER

GRADUATE SCHOOL OF SOCIAL WORK

- NGO Collaborator:

- *What do they do?* Assigning homeless youth to preventive programs against potential opioid abuse
 - Limited number of program slots, lots of homeless youth who could potentially develop opioid abuse disorders in the coming future
- *What did they need help with?* How to make the best use of the limited preventive program slots that they had?

Case Study 2: Description

- NGO Collaborator:



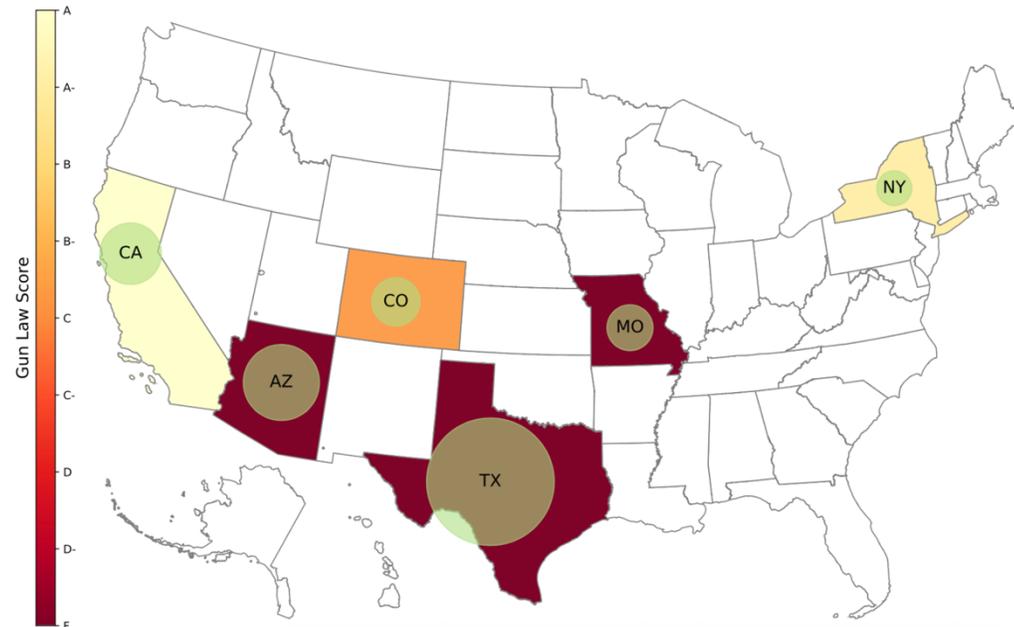
UNIVERSITY of
DENVER

GRADUATE SCHOOL OF SOCIAL WORK

- *How did we help?*

- Stage 1: Predict the likelihood of an individual youth developing opioid use disorder in the future

- Using survey data (*KDD 2020*)



Case Study 2: Description

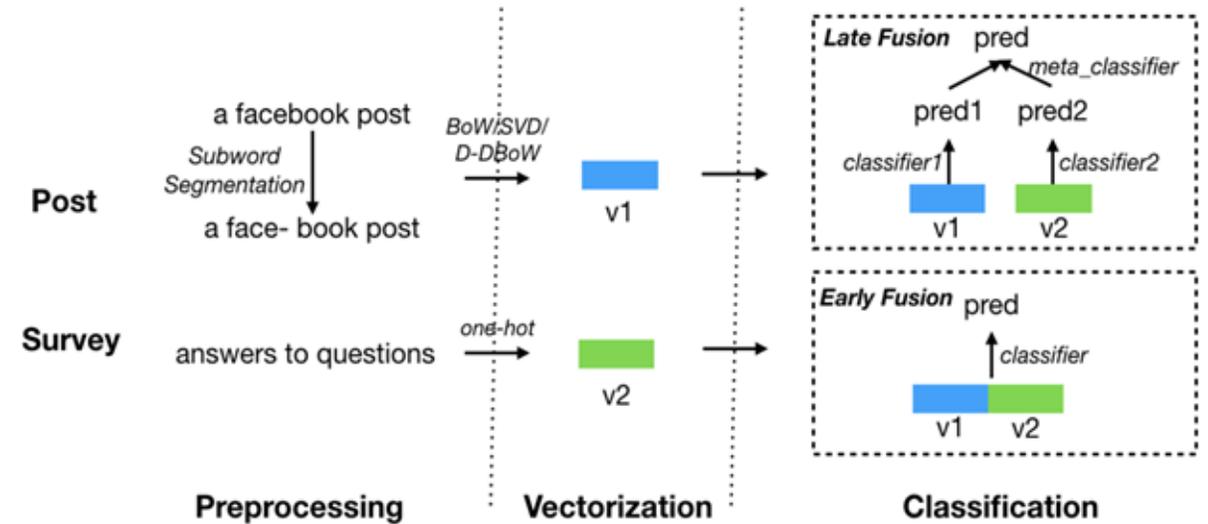
- NGO Collaborator:



- *How did we help?*

- Stage 1: Predict the likelihood of an disorder in the future

- Using survey data (*KDD 2020*)
 - Using Facebook posts (*AAAI 2021*)



“So cute!! I want one!! <3.”

Sample Quote from
Non-Substance User

“Smoke weed every day”

Sample Quote from
Substance User

Case Study 2: Description

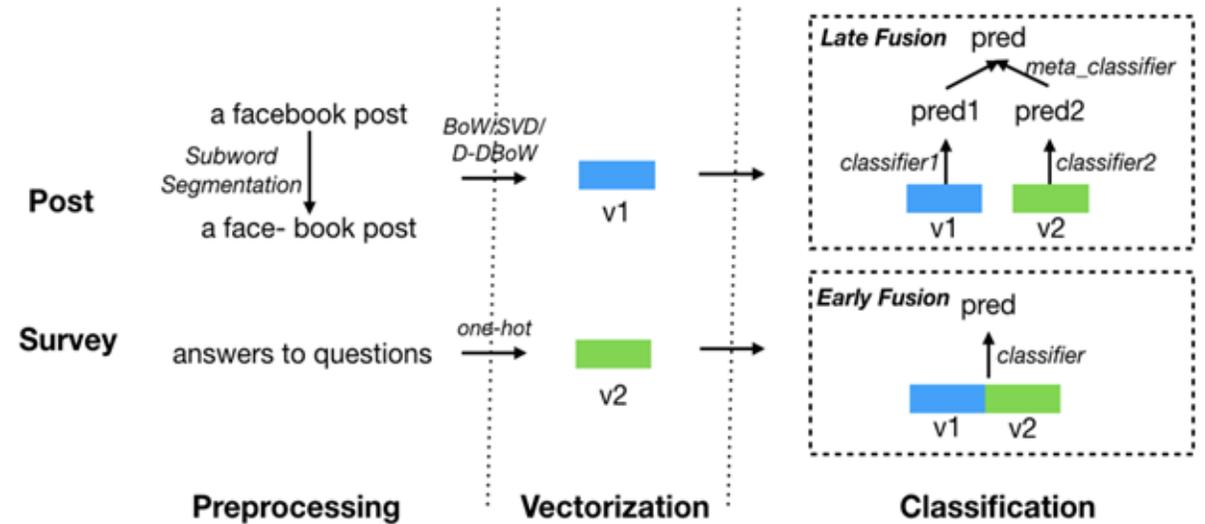
• NGO Collaborator:



• *How did we help?*

- Stage 1: Predict the likelihood of an disorder in the future

- Using survey data (*KDD 2020*)
- Using Facebook posts (*AAAI 2021*)



“So cute!! I want one!! <3.”

“Smoke weed every day”

	anger	anticipation	disgust	fear	joy	sadness	surprise	trust
Substance Users	0.419	0.505	0.325	0.346	0.544	0.331	0.252	0.594
Non-Substance Users	0.353	0.473	0.300	0.325	0.593	0.312	0.219	0.582

Case Study 2: Description

- NGO Collaborator:



UNIVERSITY of
DENVER

GRADUATE SCHOOL OF SOCIAL WORK

- *How did we help?*

- Stage 2: Use Stage 1 predictions to assign youth to program slots to maximize efficiency

- Predict-then-Optimize pipeline (*IJCAI 2020*)

$$\begin{aligned} \min_{\mathbf{x}_{ij}, \mathbf{w}_i} \quad & \sum_{i=1}^N (1 - \mathbf{w}_i) p_i^0 + \sum_{i=1}^N \sum_{j=i}^M (1 - y_{ij}) p_{ij} \mathbf{x}_{ij} \\ \text{s.t.} \quad & \sum_{j=1}^M \mathbf{x}_{ij} \leq 1, \forall i \in 1 \dots N \\ & \sum_{j=1}^M \mathbf{x}_{ij} = \mathbf{w}_i, \forall i \in 1 \dots N \\ & \sum_{i=1}^N \mathbf{x}_{ij} \leq C_j, \forall j = 1, \dots, M \\ & \mathbf{x}_{ij} \leq 1 - y_{ij}, \forall i = 1, \dots, N, j = 1, \dots, M \\ & \mathbf{x}_{ij}, \mathbf{w}_i \in \{0, 1\}, \forall i = 1, \dots, N, j = 1, \dots, M \end{aligned}$$

Case Study 3

AI for Suicide Prevention



Case Study 3: Description

- NGO Collaborator:



- *What do they do?* Gatekeeper training interventions among homeless

- Gatekeepers tasked with looking out for suicidal tendencies among their peers
 - Goal: To maximize the number of people who are cared for by a gatekeeper



Case Study 3: Description

- NGO Collaborator:



- *What did they need help with?*

Assuming partial information about a social network of youth, which people should be selected as ideal gatekeepers?



Case Study 3: Description



- NGO Collaborator: *How did we help?* Mixed Integer Bilinear program to select the optimal set of gatekeepers that could maximize network coverage (*NeurIPS 19*)
 - Some youth may say no when we enroll them to be gatekeepers
 - Want gatekeeper selection to be **robust**

$$\begin{aligned}
 & \max \quad \tau \\
 & \text{s.t.} \quad \tau \in \mathbb{R}, \mathbf{x} \in \mathcal{X}, \mathbf{y}^k \in \mathcal{Y} \quad \forall k \in \mathcal{K} \\
 & \quad \theta(\ell), \beta^k(\ell) \in \mathbb{R}_+^N, \alpha(\ell) \in \mathbb{R}_+^R, \nu(\ell) \in \mathbb{R}_+^K, \lambda(\ell) \in \Delta_K(\ell) \\
 & \quad \tau \leq -\mathbf{e}^\top \theta(\ell) + \alpha(\ell)^\top \mathbf{b} - \sum_{\substack{k \in \mathcal{K}: \\ \ell_k \neq 0}} (\mathbf{y}_{\ell_k}^k - 1) \nu_k(\ell) + \dots \\
 & \quad \dots + \sum_{\substack{k \in \mathcal{K}: \\ \ell_k = 0}} \sum_{n \in \mathcal{N}} \mathbf{y}_n^k \beta_n^k(\ell) + \sum_{k \in \mathcal{K}} \lambda_k(\ell) \sum_{n \in \mathcal{N}} \mathbf{y}_n^k \\
 & \quad \theta_n(\ell) \leq \mathbf{A}^\top \alpha(\ell) + \sum_{\substack{k \in \mathcal{K}: \\ \ell_k \neq 0}} \sum_{\nu \in \delta(\ell_k)} \mathbf{x}_\nu \nu_k(\ell) - \sum_{\substack{k \in \mathcal{K}: \\ \nu \in \delta(n)}} \sum_{\ell_k = 0} \mathbf{x}_\nu \beta_n^k(\ell) \quad \forall n \in \mathcal{N} \\
 & \quad \left. \begin{aligned}
 & \theta(\ell) \in \mathbb{R}_+^N, \alpha(\ell) \in \mathbb{R}_+^R, \nu(\ell) \in \mathbb{R}_+^K \\
 & 1 \leq -\mathbf{e}^\top \theta(\ell) + \alpha(\ell)^\top \mathbf{b} - \sum_{\substack{k \in \mathcal{K}: \\ \ell_k \neq 0}} (\mathbf{y}_{\ell_k}^k - 1) \nu_k(\ell) \\
 & \theta_n(\ell) \leq \mathbf{A}^\top \alpha(\ell) + \sum_{\substack{k \in \mathcal{K}: \\ \ell_k \neq 0}} \sum_{\nu \in \delta(\ell_k)} \mathbf{x}_\nu \nu_k(\ell) \quad \forall n \in \mathcal{N}
 \end{aligned} \right\} \forall \ell \in \mathcal{L}_+,
 \end{aligned}$$



Case Study 4

AI for Pregnant Women & New Mothers in Rural Kenya

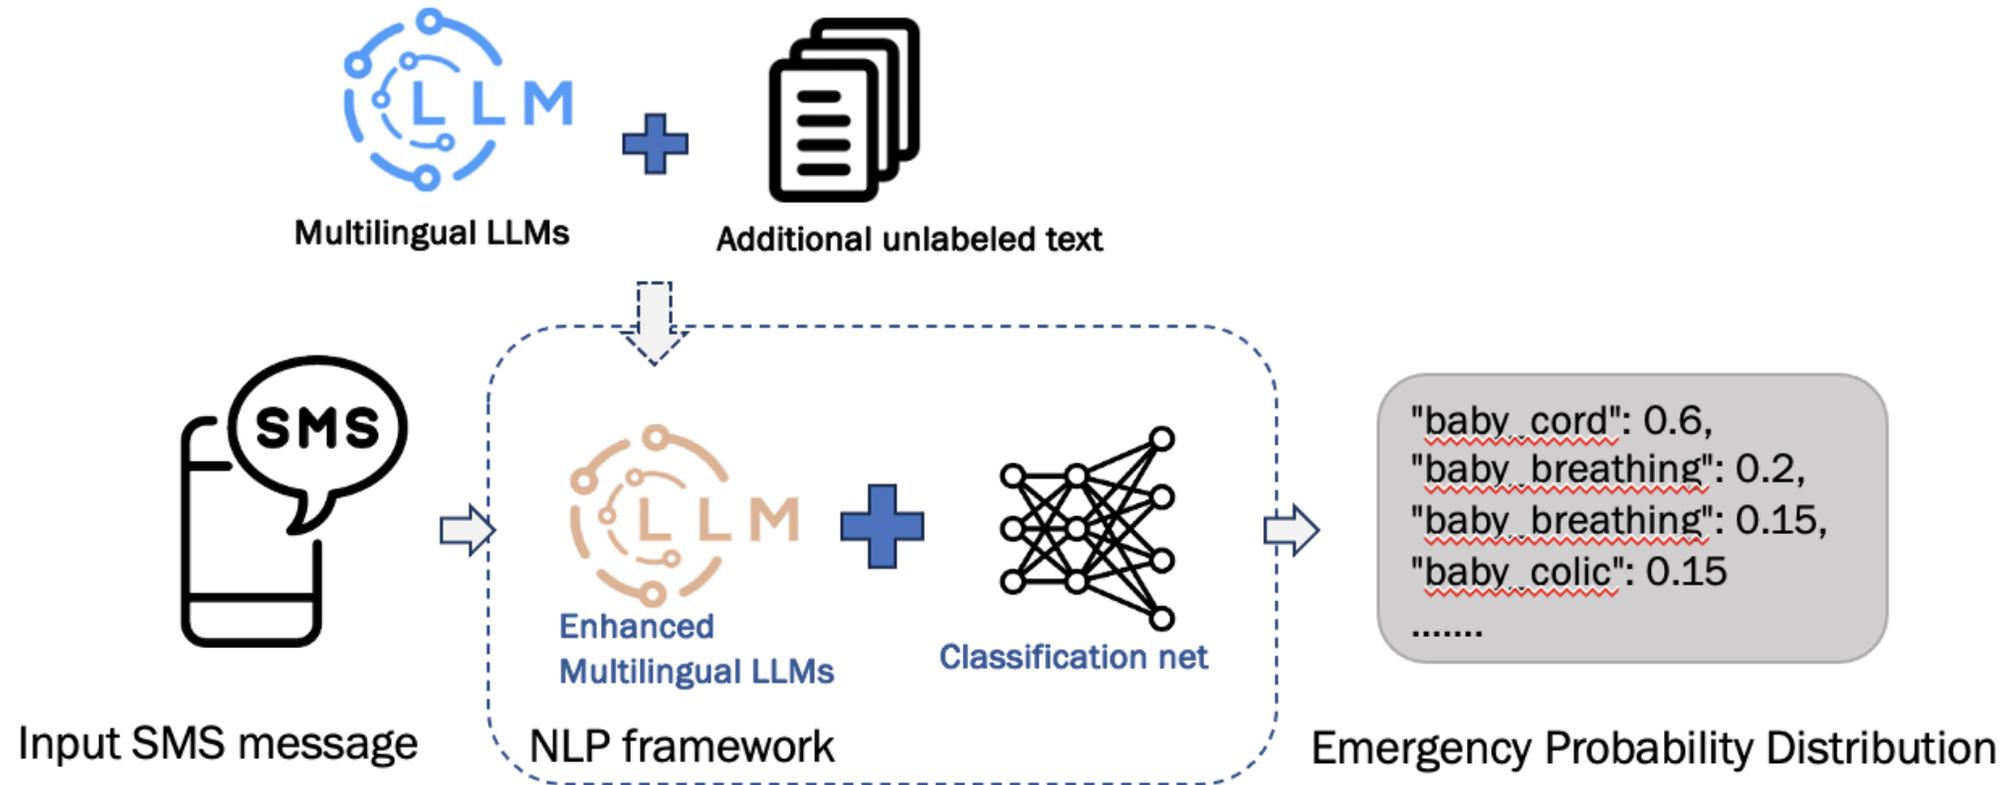
Case Study 4: Description

- NGO Collaborator:  JACARANDA HEALTH
 - *What do they do?* PROMPTS – A Tele Triage Tool to Answer Q&A for Pregnant Women
 - *What did they need help with?* All Q&A needed to be answered manually. Shortage of helpdesk operators to answer all incoming Q&A
 - *How did we help?* TRIM-AI: An AI system that reads incoming SMS messages and determines risk level, high-risk SMS should be read first

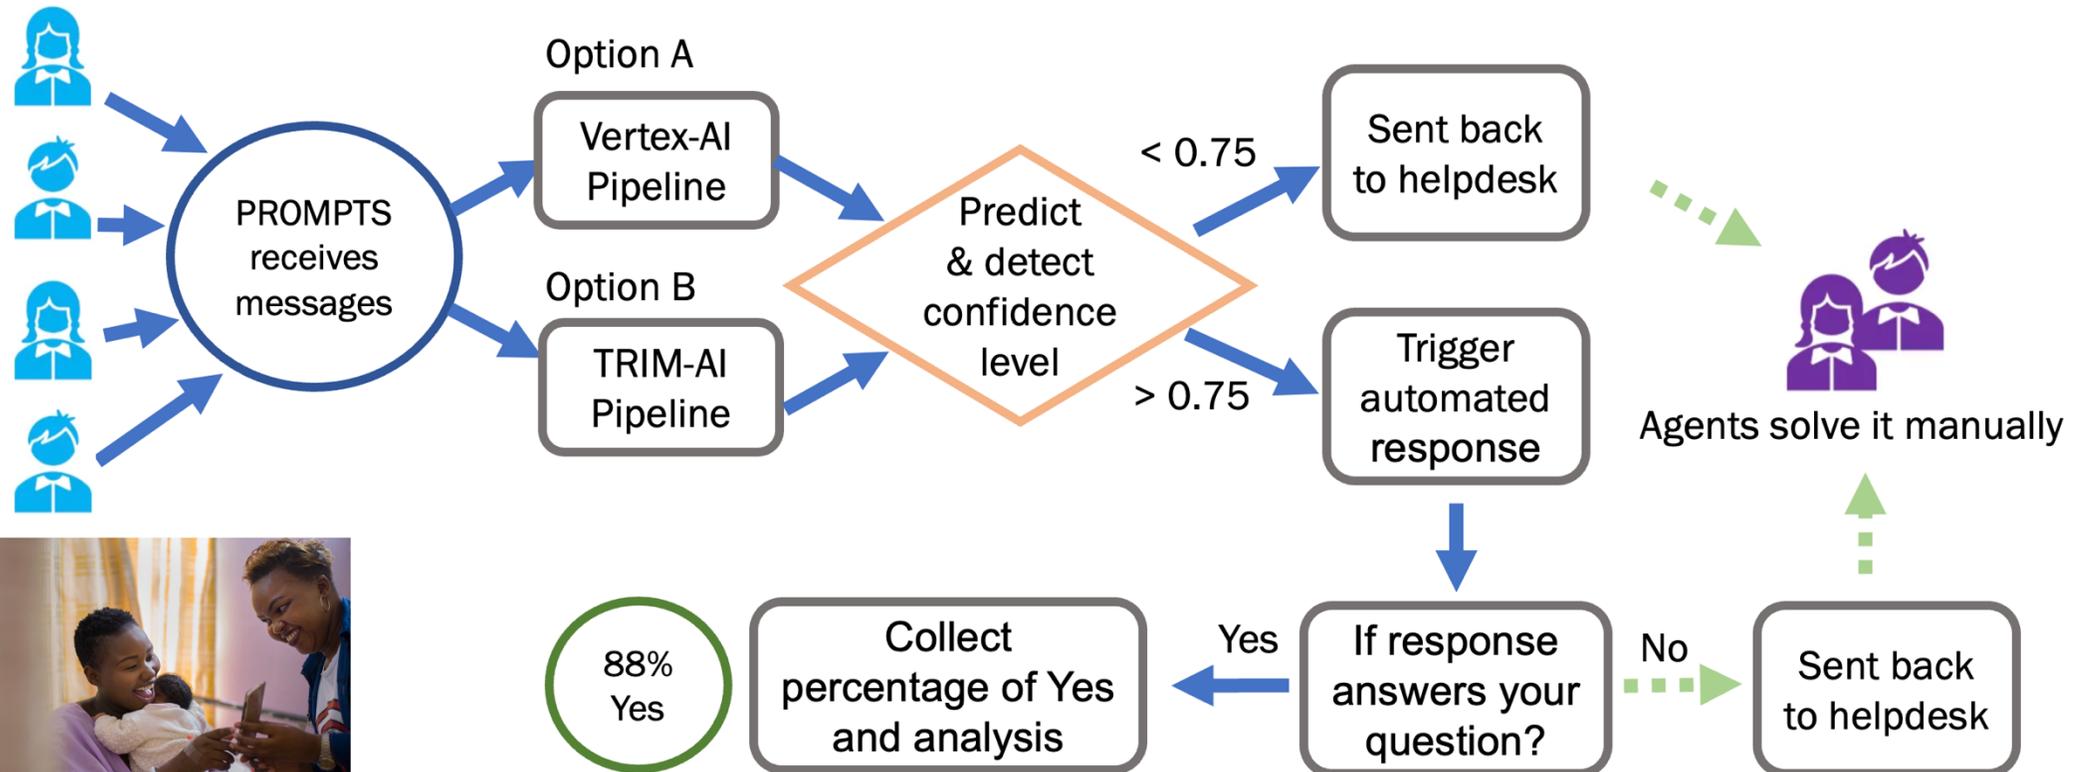
Case Study 4: High-Level Explanation



Case Study 4: High-Level Explanation

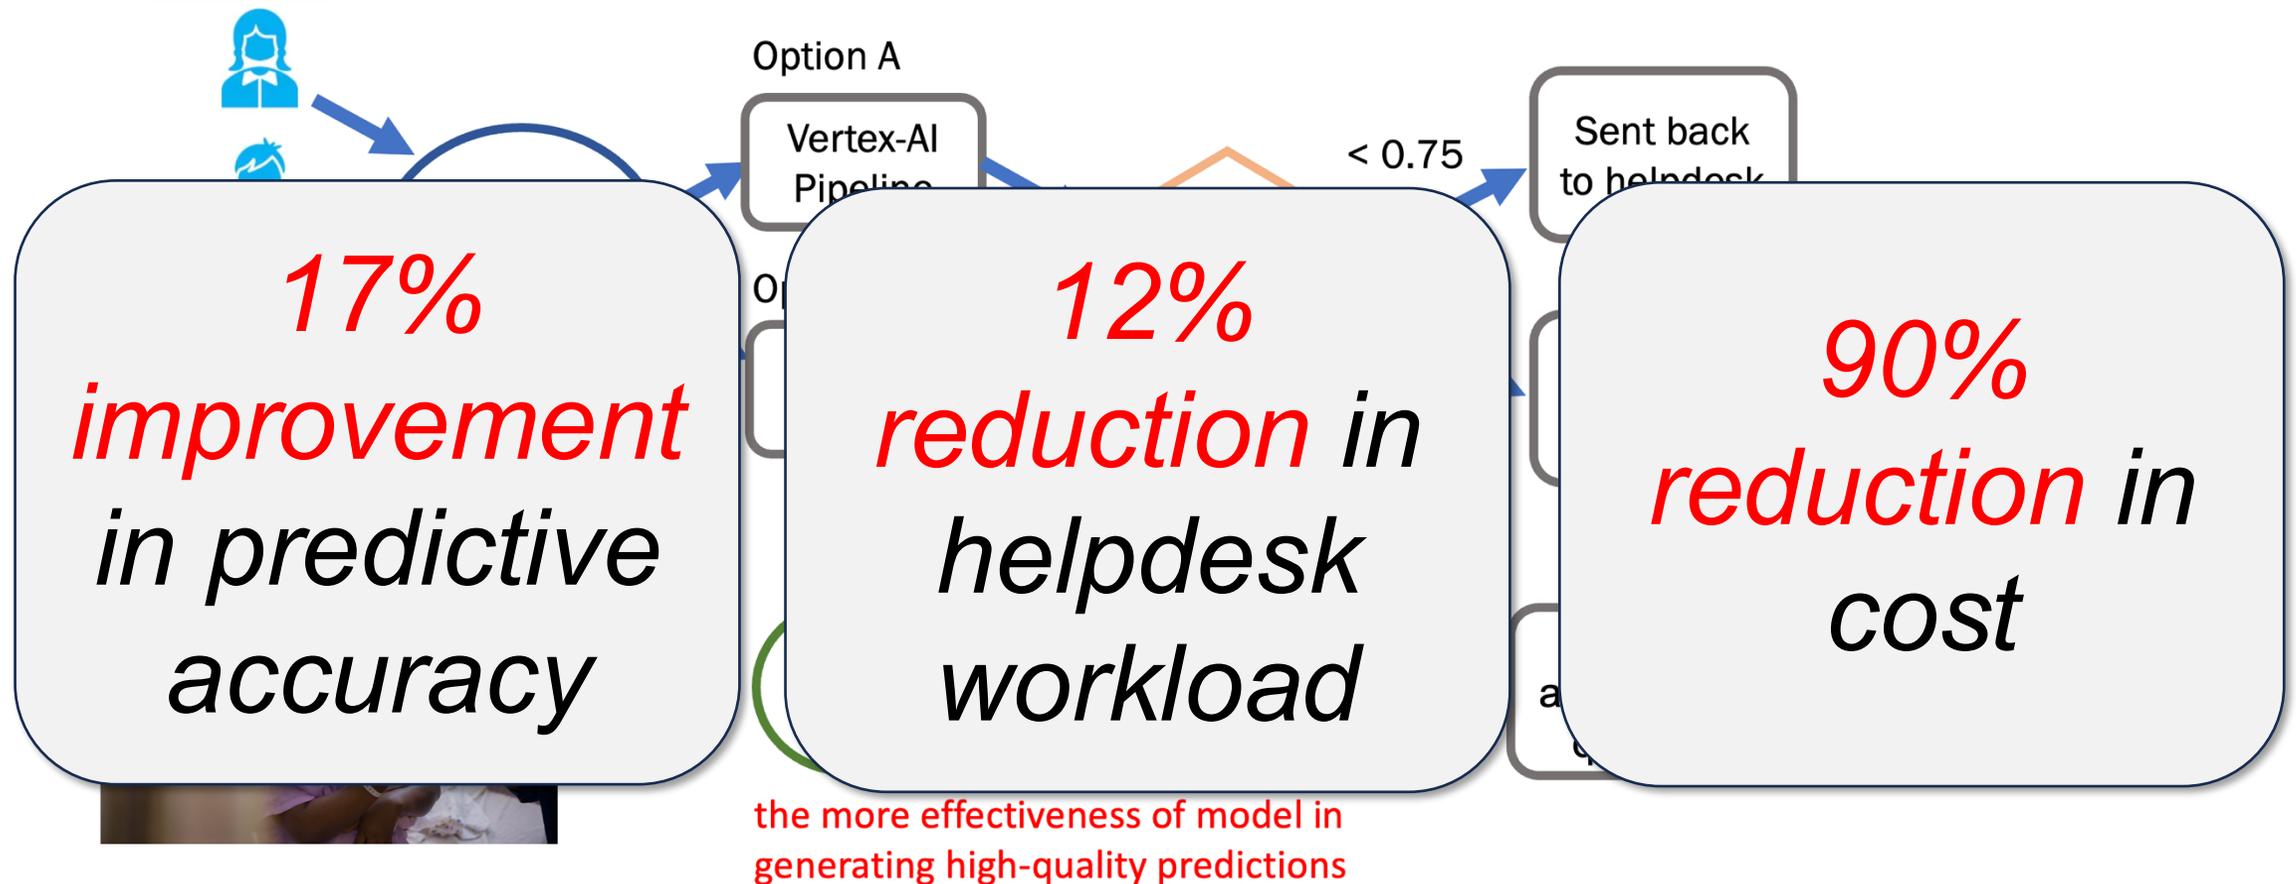


Case Study 4: Key Results



the more effectiveness of model in generating high-quality predictions

Case Study 4: Key Results



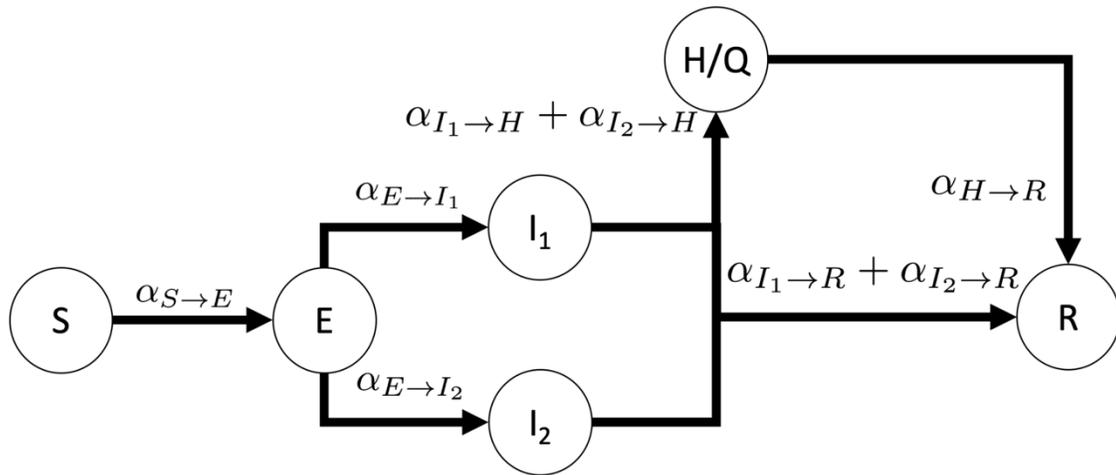


Case Study 5
POMDPs for
Testing Policy
Design for
COVID-19

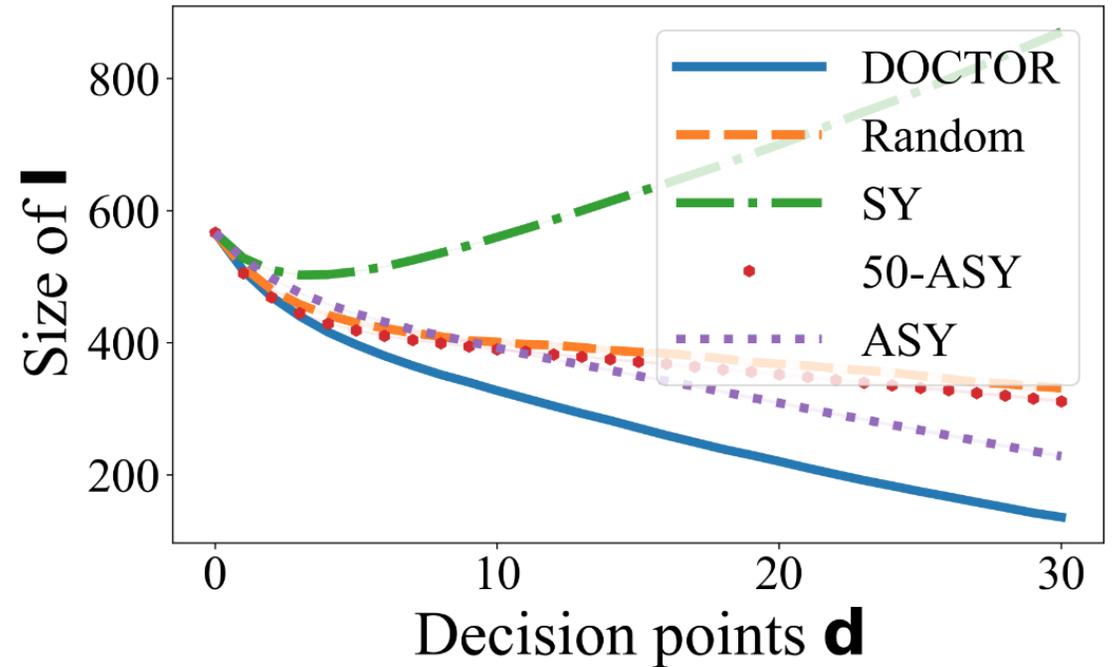
Case Study 5: Description

- NGO Collaborator:  PennState
 - *What was the problem?* COVID-19 was in full flow
 - Should Penn State open in-person for the Fall 2020 semester?
 - If so, how many COVID tests need to be procured, and how should they be allocated?
 - They were leaning towards a static solution:
 - 1200 daily tests for in-hospital testing
 - 880 daily tests (1% of student population) for random asymptomatic testing
 - *What did they need help with?* Can AI be used to develop a better testing policy
 - *How did we help?* DOCTOR: A POMDP solver that generated an optimal adaptive testing policy for determining how should next week's tests be allocated (dependent on history)

Case Study 5: Key Results



Flow Dynamics of SEIR Based Transition Model of DOCTOR's POMDP (AAMAS 2021)



DOCTOR resulted in 30% fewer infections as compared to Penn State's strategy

The gap between algorithmic success and sustainable change

Despite the technical validation, the publications, and the field tests...



2

of these systems was ever actually deployed.



2

achieved sustained impact at scale.



1

are operating today.



Deployment in AI4SG (for HIV Prevention)

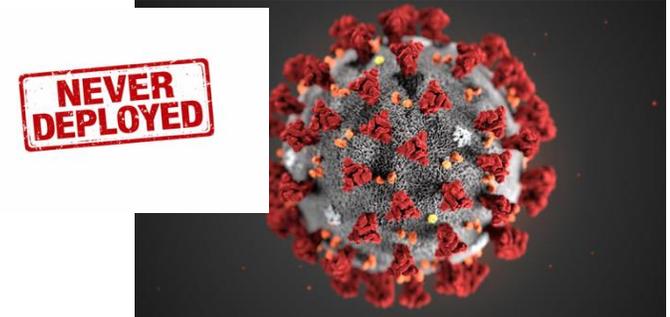
- What do successful deployments look like in AI for Social Good?
- What have they looked like in HIV prevention projects?
 - Algorithm development
 - Several pilot studies with homeless youth
 - One large multi-year clinical trial with homeless youth (ended in 2021)
 - No one is using it at the moment 😞

Deployment in AI4SG

AI for Substance Abuse

AI for Suicide Prevention

AI for POMDP Testing



Cause of Failure

No Buy-in from NGO Leadership

Cause of Failure

Dynamic & Divergent Priorities

Cause of Failure

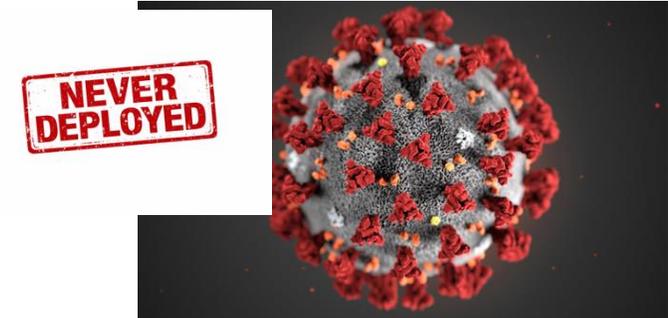
Territorial Disputes

Deployment in AI4SG (in other domains)

AI for Substance Abuse

AI for Suicide Prevention

AI for POMDP Testing



Cause of Failure

No Buy-in from NGO Leadership

Cause of Failure

Dynamic & Divergent Priorities

Cause of Failure

Territorial Disputes

AI for Prevention of Online Child Sexual Abuse



Cause of Failure

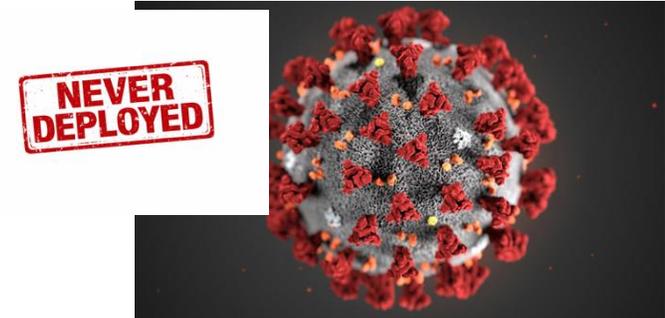
Improper Perceptions of AI

Deployment in AI4SG (in other domains)

AI for Substance Abuse

AI for Suicide Prevention

AI for POMDP Testing



Cause of Failure

No Buy-in from NGO Leadership

Cause of Failure

Dynamic & Divergent Priorities

Cause of Failure

Territorial Disputes

AI for Prevention of Online Child Sexual Abuse

AI for Mental Health Triage



Cause of Failure

Improper Perceptions of AI

Cause of Failure

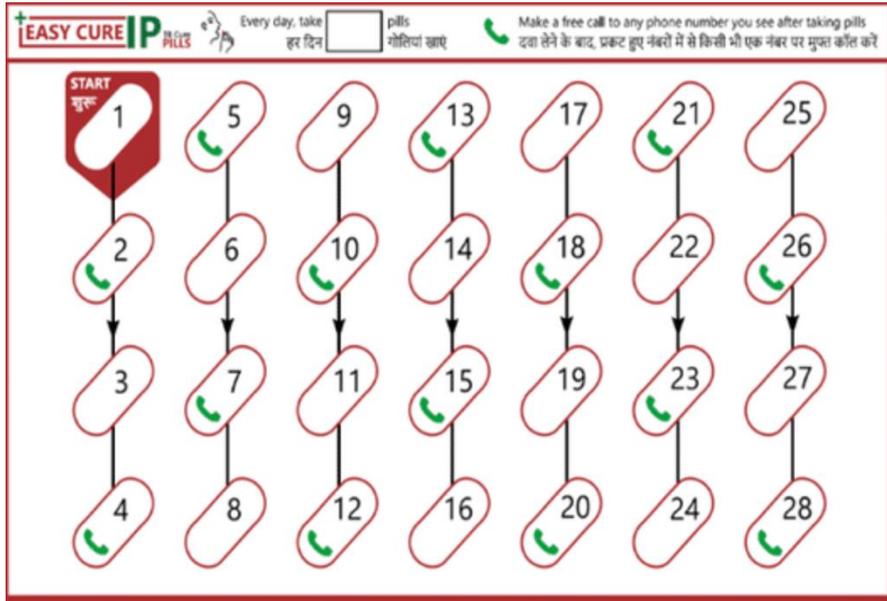
Red Tape in Getting Data Agreements

Different Failure Modes - AI4SG Deployment



Majumdar et al., The Hardness of Achieving Impact in AI for Social Impact Research (FaCCT 26 Submission)

A Contrastive Case: 99DOTS



Last Dosage

- Bhramdeov P**
Last Dosage: 3 days ago
- Dablu na**
Last Dosage: 3 days ago
- Gyanti kumar**
Last Dosage: 3 days ago
- Saheen kumar**
Last Dosage: 2 days ago
- Gautam Soni**
Last Dosage: 2 days ago
- Miss ranjan**
Last Dosage: 2 days ago
- baalak r**

Touch for summary
% Called Today: 80%

(a)

Patient Details

Prem Kumar [Call](#)

MARCH 2018

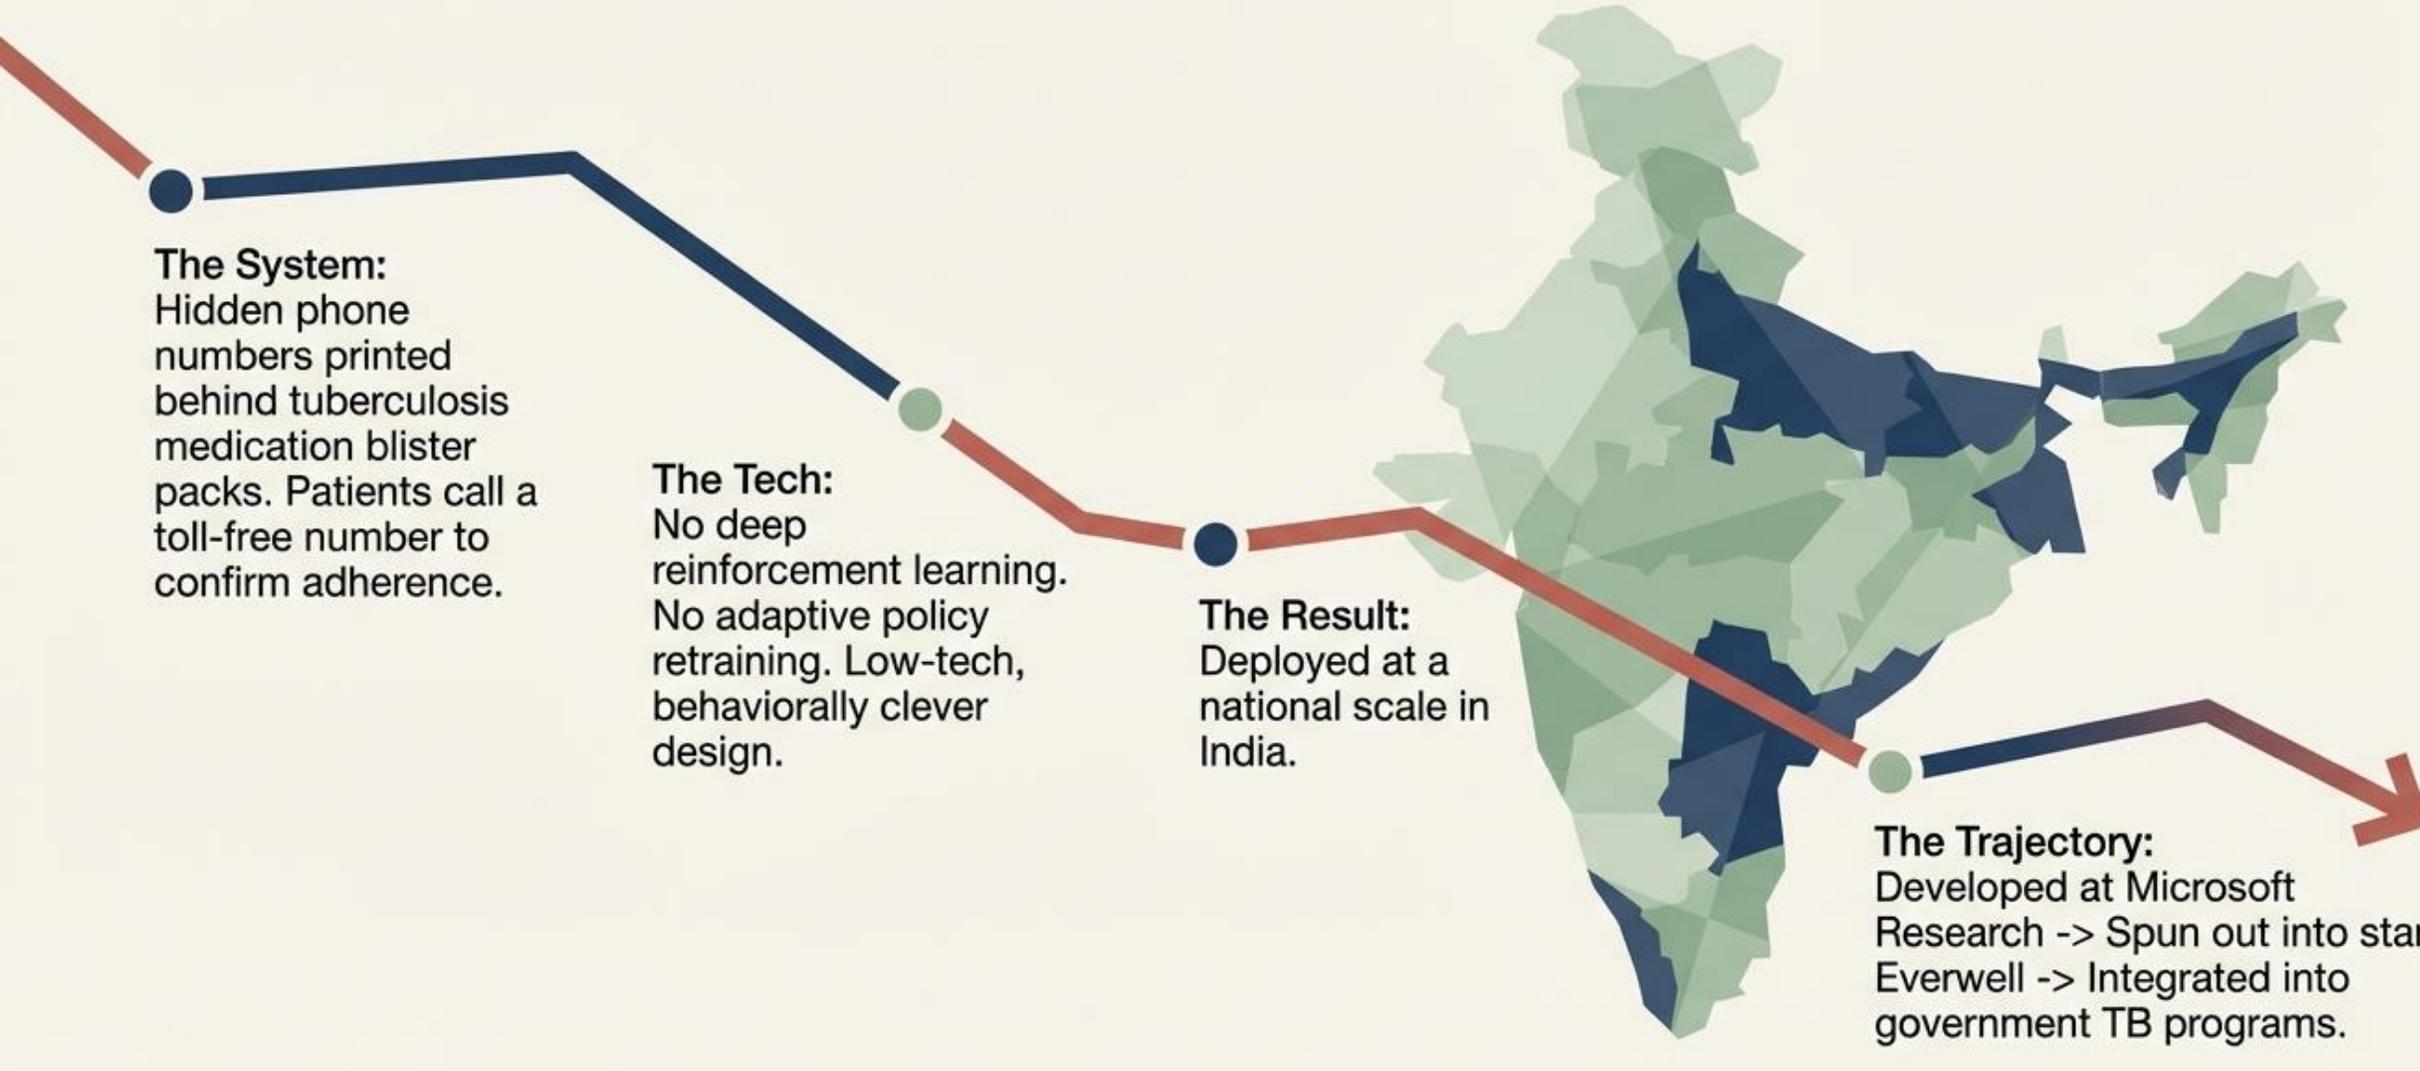
SUN	MON	TUE	WED	THU	FRI	SAT
25	26	27	28	1	2	3
4	5	6	7	8	9	10
11	12	13	14	15	16	17
18	19	20	21	22	23	24
25	26	27	28	29	30	31

Manual Doses: [Add](#) [Remove](#)

Basic Medical

(b)

99DOTS optimized for integration, not intelligence



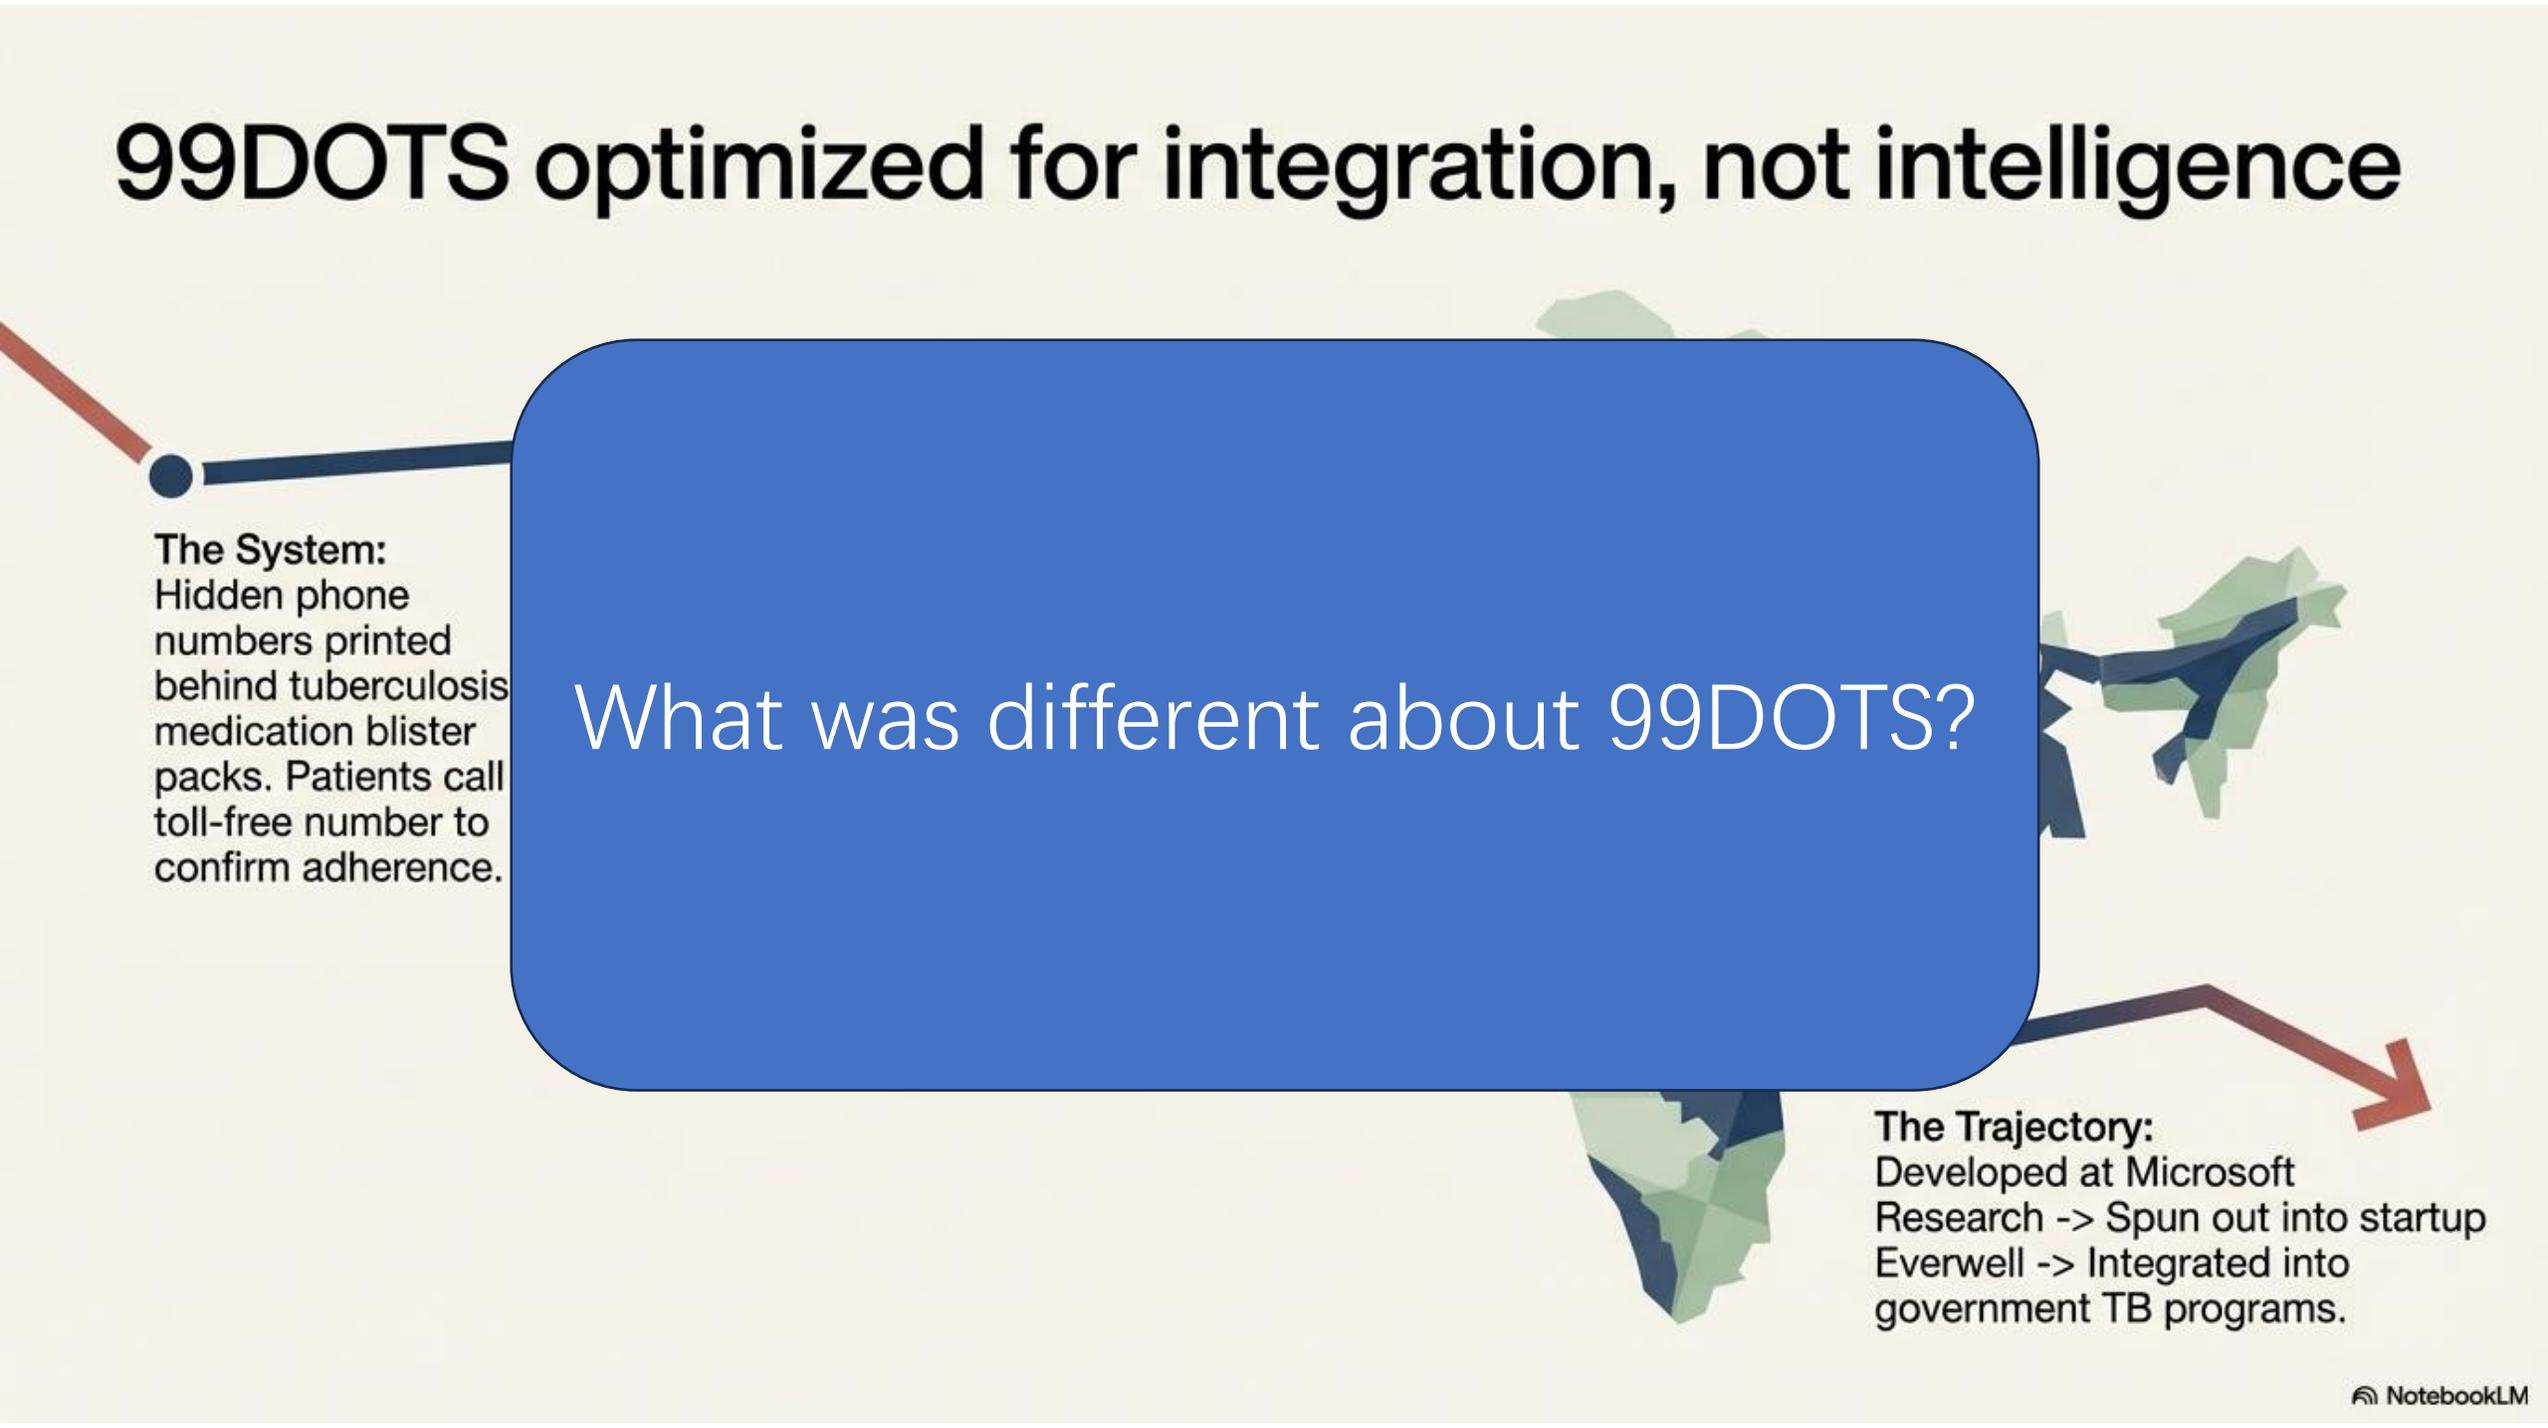
The System:
Hidden phone numbers printed behind tuberculosis medication blister packs. Patients call a toll-free number to confirm adherence.

The Tech:
No deep reinforcement learning. No adaptive policy retraining. Low-tech, behaviorally clever design.

The Result:
Deployed at a national scale in India.

The Trajectory:
Developed at Microsoft Research -> Spun out into startup Everwell -> Integrated into government TB programs.

99DOTS optimized for integration, not intelligence



The System:
Hidden phone numbers printed behind tuberculosis medication blister packs. Patients call toll-free number to confirm adherence.

What was different about 99DOTS?

The Trajectory:
Developed at Microsoft Research -> Spun out into startup Everwell -> Integrated into government TB programs.

Why did 99DOTS live beyond the research stage? – Technical Simplicity

- Was it because of its technical simplicity?
 - Easier to explain to stakeholders, no sense of "black-box AI magic"
- Conversely, were AI4SG projects doomed for "failure" because we optimize for acceptance into AI conferences?
 - Need fancy math, elegant theorems to get into these conferences

Why did 99DOTS live beyond the research stage? – Technical Simplicity

- Why is it hard to explain what this is doing to govt stakeholders
- Why is it hard to optimize for acceptance into the system

Hard to explain what this is doing to govt stakeholders

Sometimes, “complex” solutions may be the only way to solve the stated problem

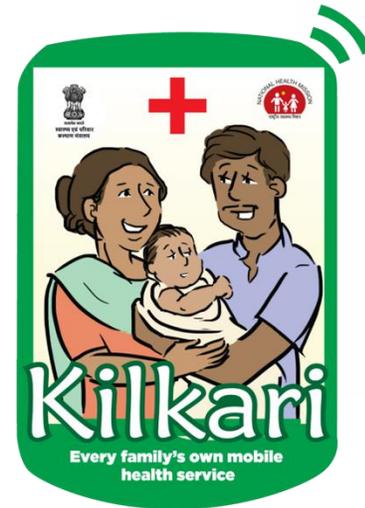
$$\begin{aligned}
 & \max \quad \tau \\
 & \text{s.t.} \quad \tau \in \mathbb{R}, \mathbf{x} \in \mathcal{X}, \mathbf{y}^k \in \mathcal{Y} \quad \forall k \in \mathcal{K} \\
 & \quad \boldsymbol{\theta}(\boldsymbol{\ell}), \boldsymbol{\beta}^k(\boldsymbol{\ell}) \in \mathbb{R}_+^N, \boldsymbol{\alpha}(\boldsymbol{\ell}) \in \mathbb{R}_+^R, \boldsymbol{\nu}(\boldsymbol{\ell}) \in \mathbb{R}_+^K, \boldsymbol{\lambda}(\boldsymbol{\ell}) \in \Delta_K(\boldsymbol{\ell}) \\
 & \quad \tau \leq -\mathbf{e}^\top \boldsymbol{\theta}(\boldsymbol{\ell}) + \boldsymbol{\alpha}(\boldsymbol{\ell})^\top \mathbf{b} - \sum_{\substack{k \in \mathcal{K}: \\ \ell_k \neq 0}} (\mathbf{y}_{\ell_k}^k - 1) \boldsymbol{\nu}_k(\boldsymbol{\ell}) + \dots \\
 & \quad \dots + \sum_{\substack{k \in \mathcal{K}: \\ \ell_k = 0}} \sum_{n \in \mathcal{N}} \mathbf{y}_n^k \boldsymbol{\beta}_n^k(\boldsymbol{\ell}) + \sum_{k \in \mathcal{K}} \lambda_k(\boldsymbol{\ell}) \sum_{n \in \mathcal{N}} \mathbf{y}_n^k \\
 & \quad \boldsymbol{\theta}_n(\boldsymbol{\ell}) \leq \mathbf{A}^\top \boldsymbol{\alpha}(\boldsymbol{\ell}) + \sum_{\substack{k \in \mathcal{K}: \\ \ell_k \neq 0}} \sum_{\nu \in \delta(\ell_k)} \mathbf{x}_\nu \boldsymbol{\nu}_k(\boldsymbol{\ell}) - \sum_{\substack{k \in \mathcal{K}: \\ \ell_k = 0}} \sum_{\nu \in \delta(n)} \mathbf{x}_\nu \boldsymbol{\beta}_n^k(\boldsymbol{\ell}) \quad \forall n \in \mathcal{N} \quad \forall n \in \mathcal{N} \\
 & \quad \left. \begin{aligned}
 & \boldsymbol{\theta}(\boldsymbol{\ell}) \in \mathbb{R}_+^N, \boldsymbol{\alpha}(\boldsymbol{\ell}) \in \mathbb{R}_+^R, \boldsymbol{\nu}(\boldsymbol{\ell}) \in \mathbb{R}_+^K \\
 & 1 \leq -\mathbf{e}^\top \boldsymbol{\theta}(\boldsymbol{\ell}) + \boldsymbol{\alpha}(\boldsymbol{\ell})^\top \mathbf{b} - \sum_{\substack{k \in \mathcal{K}: \\ \ell_k \neq 0}} (\mathbf{y}_{\ell_k}^k - 1) \boldsymbol{\nu}_k(\boldsymbol{\ell}) \\
 & \boldsymbol{\theta}_n(\boldsymbol{\ell}) \leq \mathbf{A}^\top \boldsymbol{\alpha}(\boldsymbol{\ell}) + \sum_{\substack{k \in \mathcal{K}: \\ \ell_k \neq 0}} \sum_{\nu \in \delta(\ell_k)} \mathbf{x}_\nu \boldsymbol{\nu}_k(\boldsymbol{\ell}) \quad \forall n \in \mathcal{N}
 \end{aligned} \right\} \forall \boldsymbol{\ell} \in \mathcal{L}_+,
 \end{aligned}$$

Why did 99DOTS live beyond the research stage? – Institutional Support

- Was it because 99DOTS was developed within the Microsoft ecosystem?
 - Easy access to well-trained software development teams
 - Makes it easier to develop POCs
- Conversely, who does this work in academic AI4SG projects?
 - PhD students, who are not being trained for software dev
 - Maybe undergrad students, but steep learning curve

Why did 99DOTS live beyond the research stage? – Local Capacity

- Was it because there was local capacity available which could “own” the project sustainably?



No Funding for
Tech Initiatives

Dedicated Tech Teams

Why did 99DOTS live beyond the research stage? – Minimal Effort Required to Integrate

- Was it because the data requirements were minimal?
 - Compared to AI4SG projects which often do not start until data is received from the non-profit side

AI for Mental Health Triage

FINGERS
CROSSED



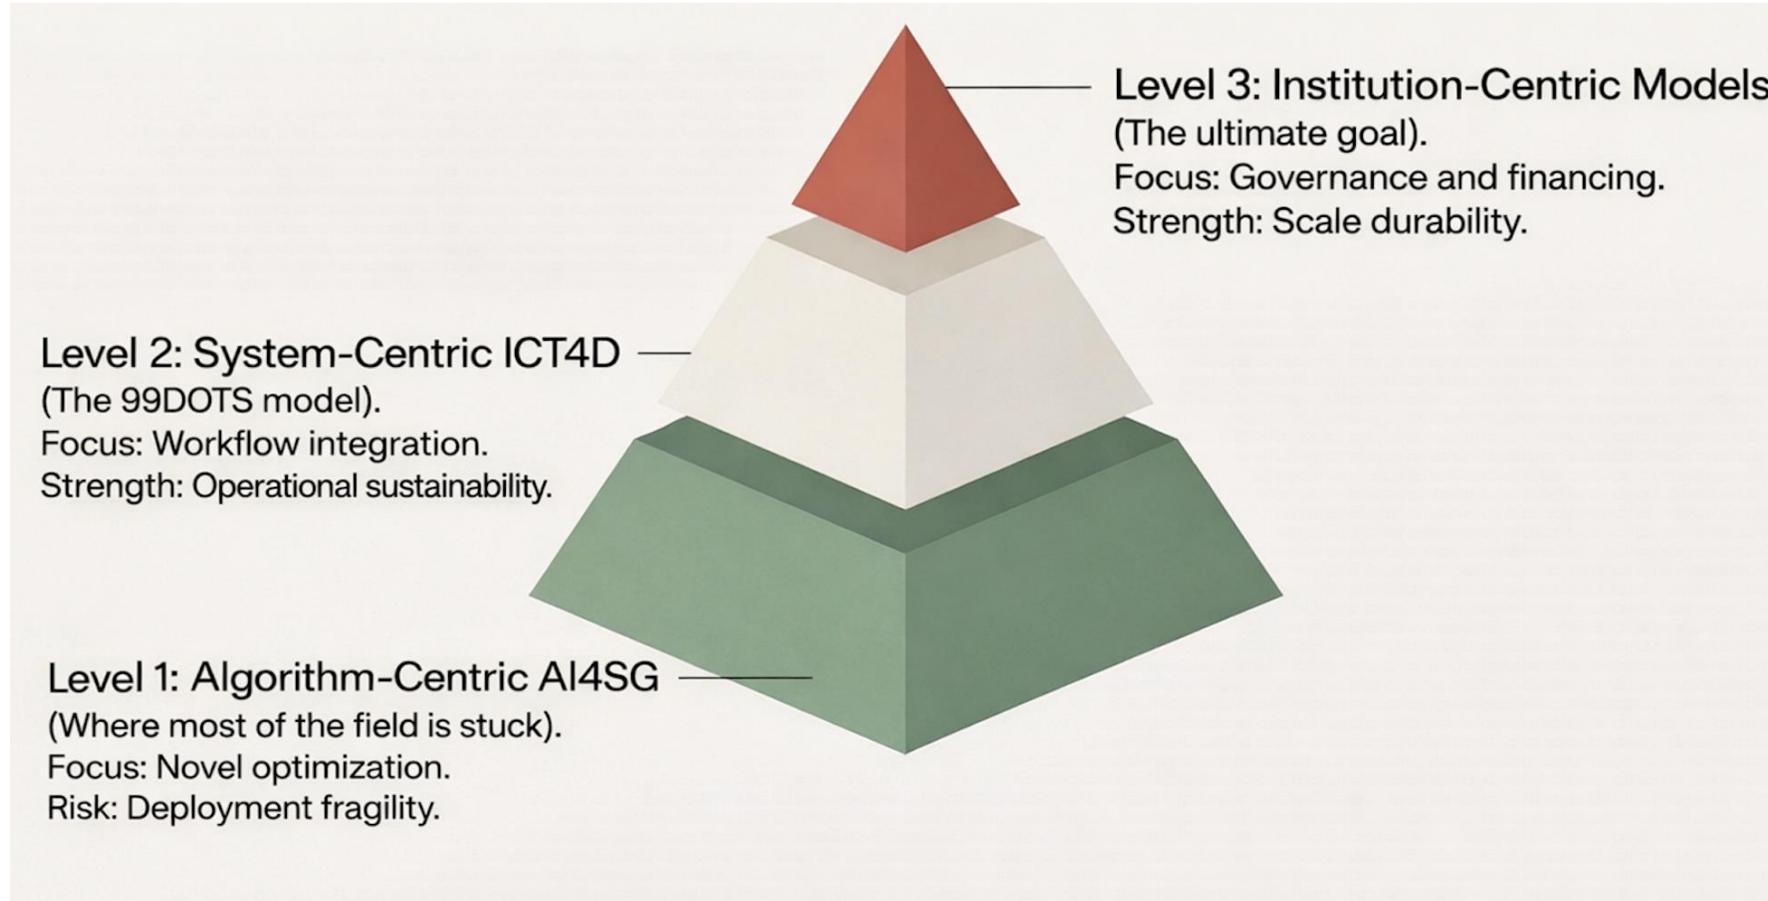
Cause of Failure

Red Tape in Getting Data Agreements

What can we do about these issues?

- Do we even want to?
 - Is it even our job as academics to achieve sustainable change?
- Should we be more selective in how we decide problems to work on?
 - Simple Problems? – At odds with current academic structures
 - Simple Solutions? - At odds with current academic structures
 - Partner with well-funded non-profits who have lots of funding? – Seems exclusionary

A Maturity Model for AI4SG Research



A Maturity Model for AI4SG Research

