MoodScope:

Sensing mood from smartphone usage patterns

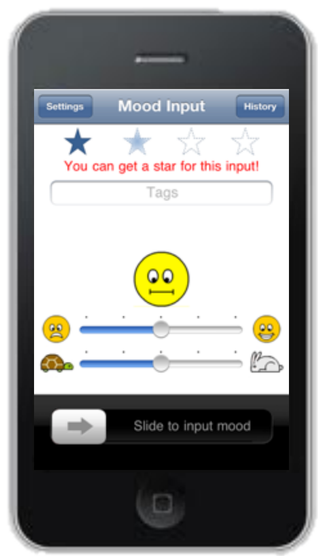

RICE

Robert Likamwa
Lin Zhong

Microsoft Research Asia

Yunxin Liu
Nicholas D. Lane

Mood-Enhanced Apps

Personal analytics

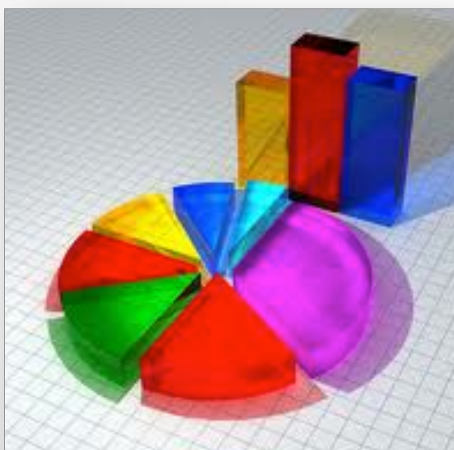

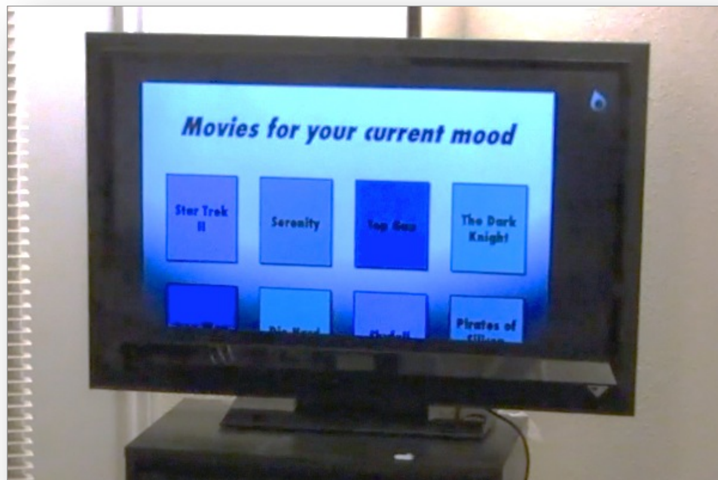

Media recommendation

Social ecosystems

Affective Computing
(Mood and Emotion)

Audio/Video-based
(AffectAura, EmotionSense)

Captures expressions
Power hungry
Slightly invasive

Biometric-based
(Skin conductivity, Temperature, Pulse rate)

Highly temporal
High cost of deployment
Hassle

Can your mobile phone infer your mood?

From already-available, low-power information?*

* No audio/video sensing, no body-instrumentation

MoodScope ∈ Affective Computing

Usage Trace-based
(MoodScope)

Passive, Continuous

How to model mood?

Audio/Video-based

Captures expressions

Power hungry

Slightly invasive

Biometric-based

Very direct, Fine-grained

High cost of deployment

Mood is…

- … a persistent long-lasting state
  - Lasts hours or days
  - Emotion lasts seconds or minutes
- … a strong social signal
  - Drives communications
  - Drives interactions
  - Drives activity patterns

Circumplex model (Russell 1980)

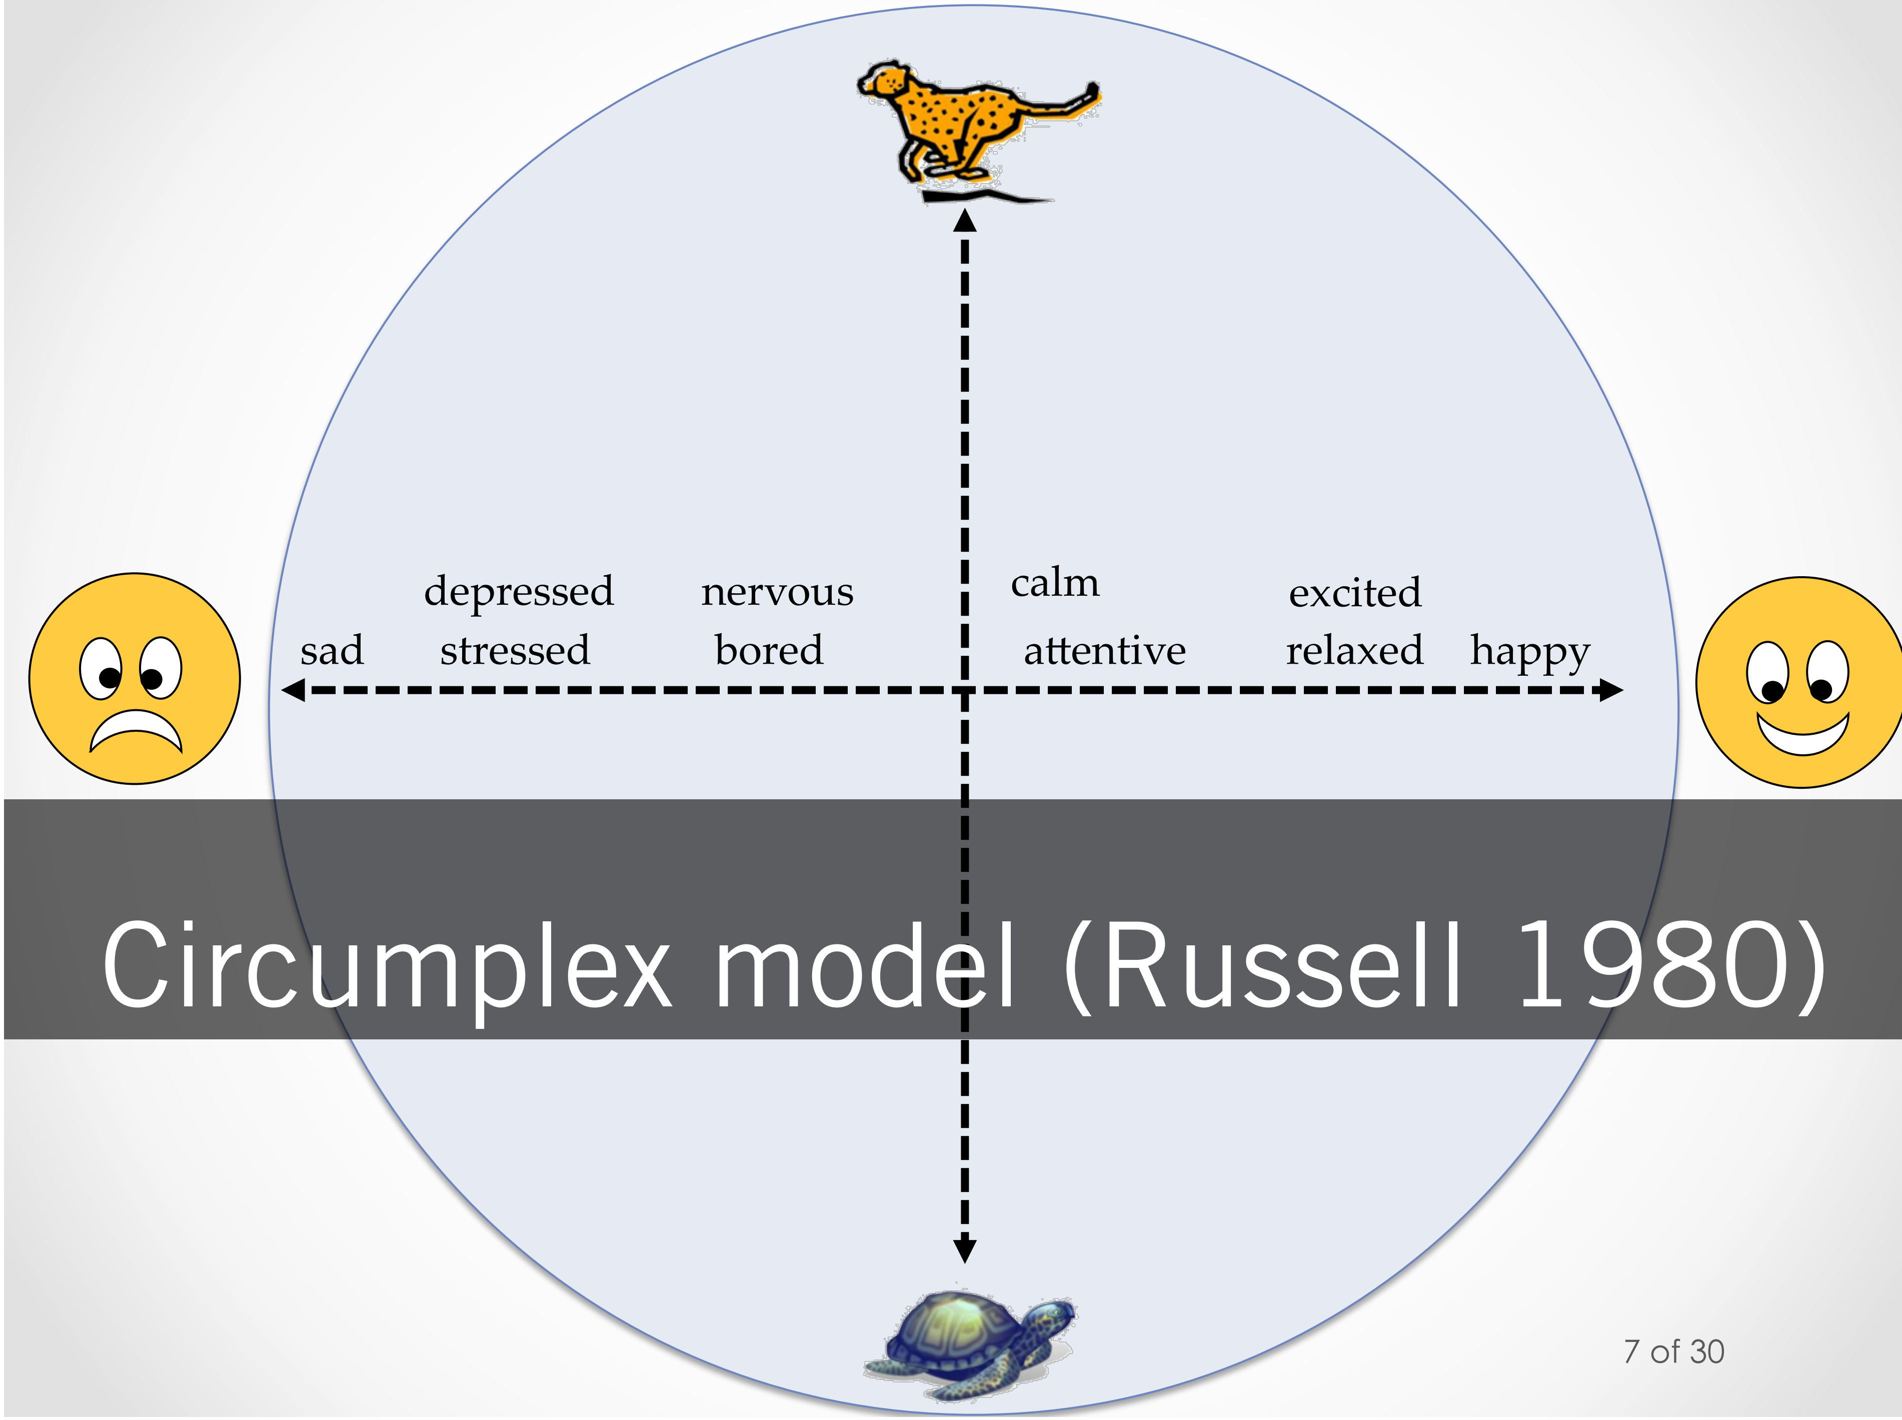

How is the user communicating?

What apps is the user using?

f (

) = *mood*

iPhone Livelab Logger

- Web history
- Phone call history
- Sms history
- Email history
- Location history
- App usage

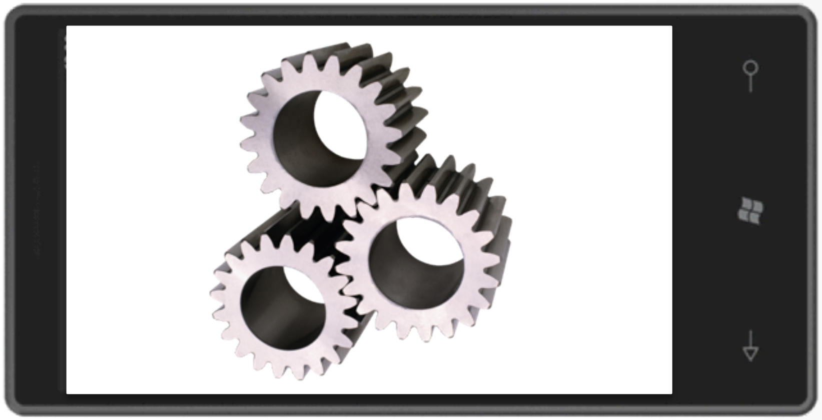

Runs as shell
Hash private data
Nightly uploads

Adapted From C. Shepard, A. Rahmati, C. Tossel, L. Zhong, And P. Kortum, "Livelab: Measuring Wireless Networks And Smartphone Users In The Field," In *Hotmetrics*, 2010.

Mood Journaling App

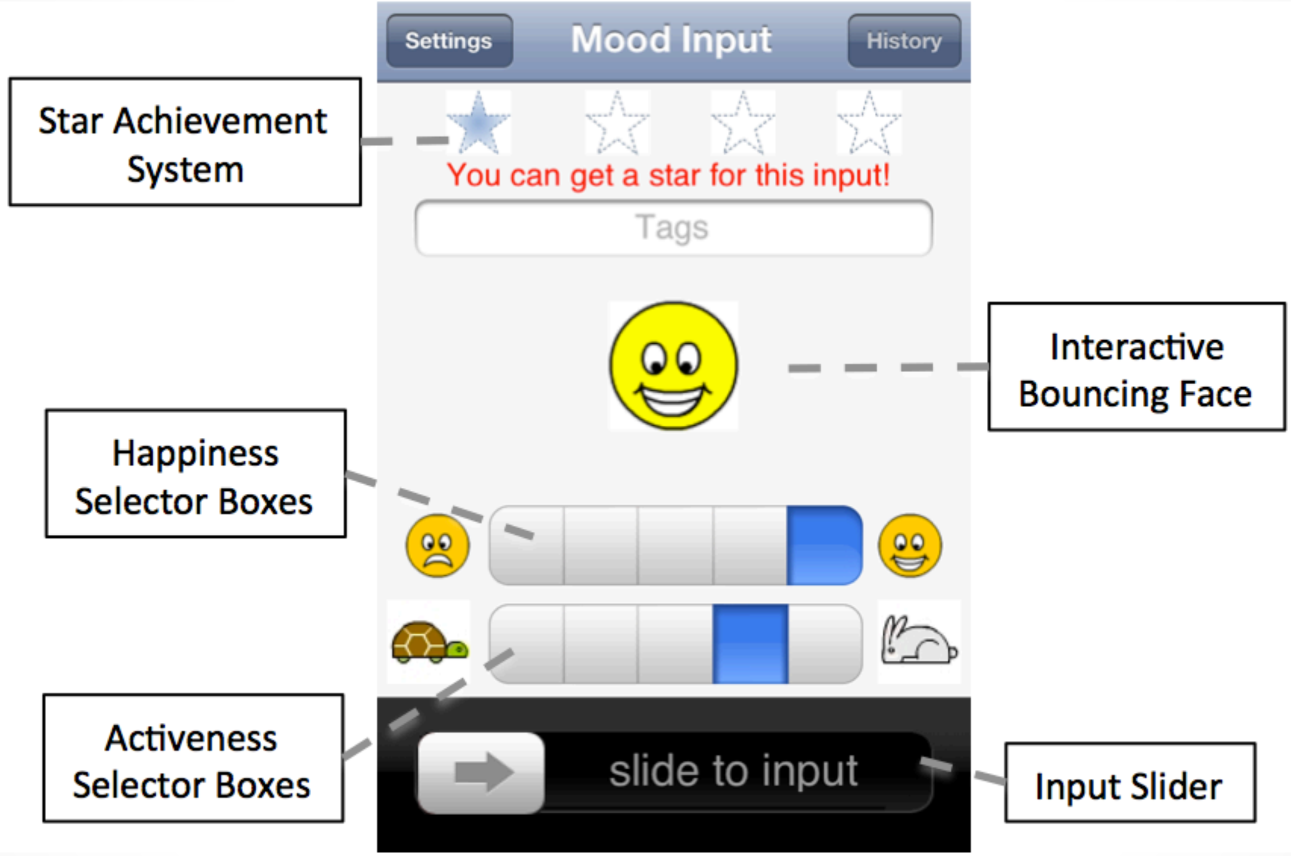

<u>User-base</u>

32 users aged between 18 and 29

11 females

Inference

- Detect a mood pattern
- Validate with only 60 days of data
- Wide range of candidate usage data
- Low computational resources

Daily Mood Averages

- Separate pleasure, activeness dimension

- Take the average over a day

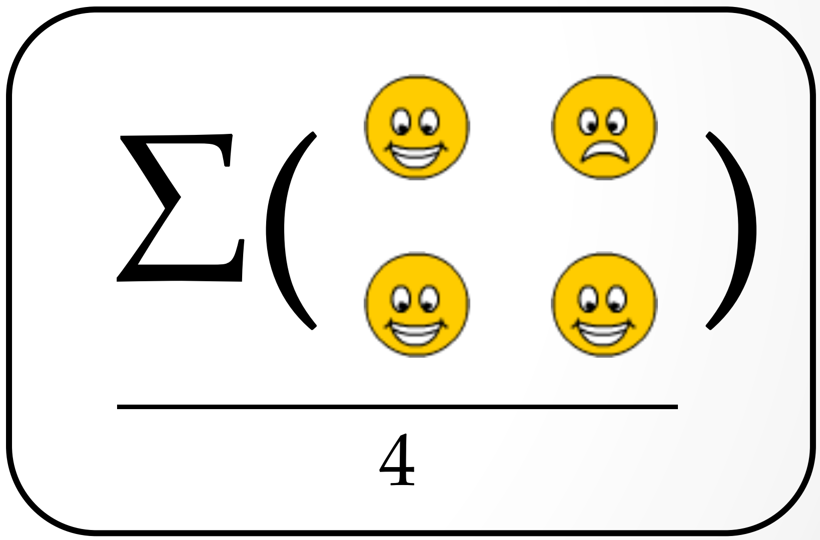

Exploring Features

- Communication
  - SMS
  - Email
  - Phone Calls

- To whom?
  - # messages
  - Length/Duration

Consider "Top 10" Histograms

? How **many** phone calls were made to #1? #2? … #10?

? How **much time** was spent on calls to #1? #2? … #10?

Exploring Features

- Communication
  - SMS
  - Email
  - Phone Calls
- To whom?
  - # messages
  - Length/Duration
- Usage Activity
  - Applications
  - Websites visited
  - Location History
- Which (app/site/location)?
  - # instances

Previous Mood

- Use previous 2 pairs of mood labels

Data Type	Histogram by:	Dimensions
Email contacts	# Messages	10
	# Characters	10
SMS contacts	# Messages	10
	# Characters	10
Phone call contacts	# Calls	10
	Call Duration	10
Website domains	# Visits	10
Location Clusters	# Visits	10
Apps	# App launches	10
	App Duration	10
Categories of Apps	# App launches	12
	App Duration	12
Previous Pleasure and Activeness Averages	N/A	4

Data Type	Histogram by:	Dimensions
Email contacts	# Messages	10
	# Characters	10
SMS contacts	# Messages	10
	# Characters	10
Phone call contacts	# Calls	10
	Call Duration	10
Website domains	# Visits	10
Location Clusters	# Visits	10
Apps	# App launches	10
	App Duration	10
Categories of Apps	# App launches	12
	App Duration	12
Previous Pleasure and Activeness Averages	N/A	4

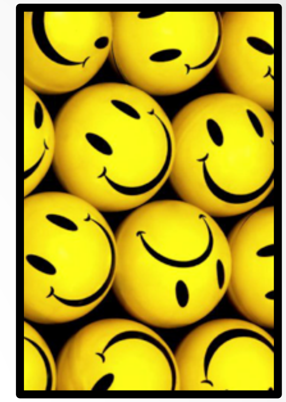

Model Design

- Multi-Linear Regression
  - Minimize Mean Squared Error
- Leave-One-Out Cross-Validation
- Sequential Forward Feature Selection during training

Sequential Feature Selection

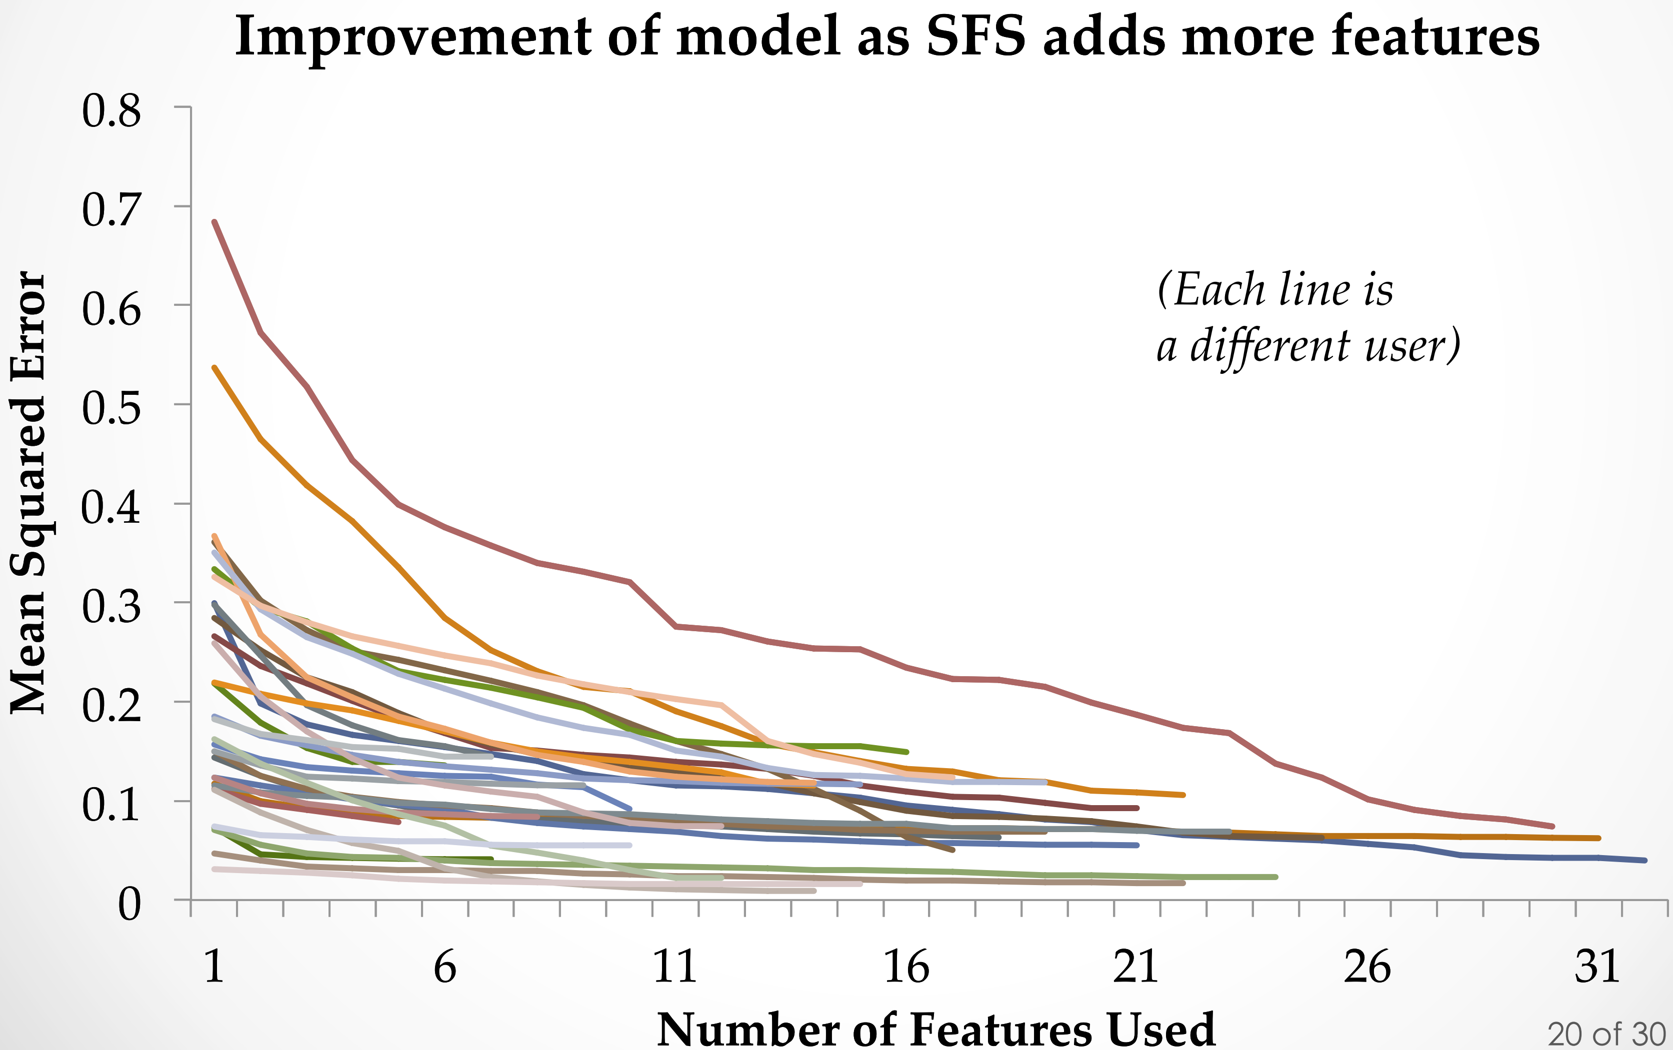

Sample Prediction

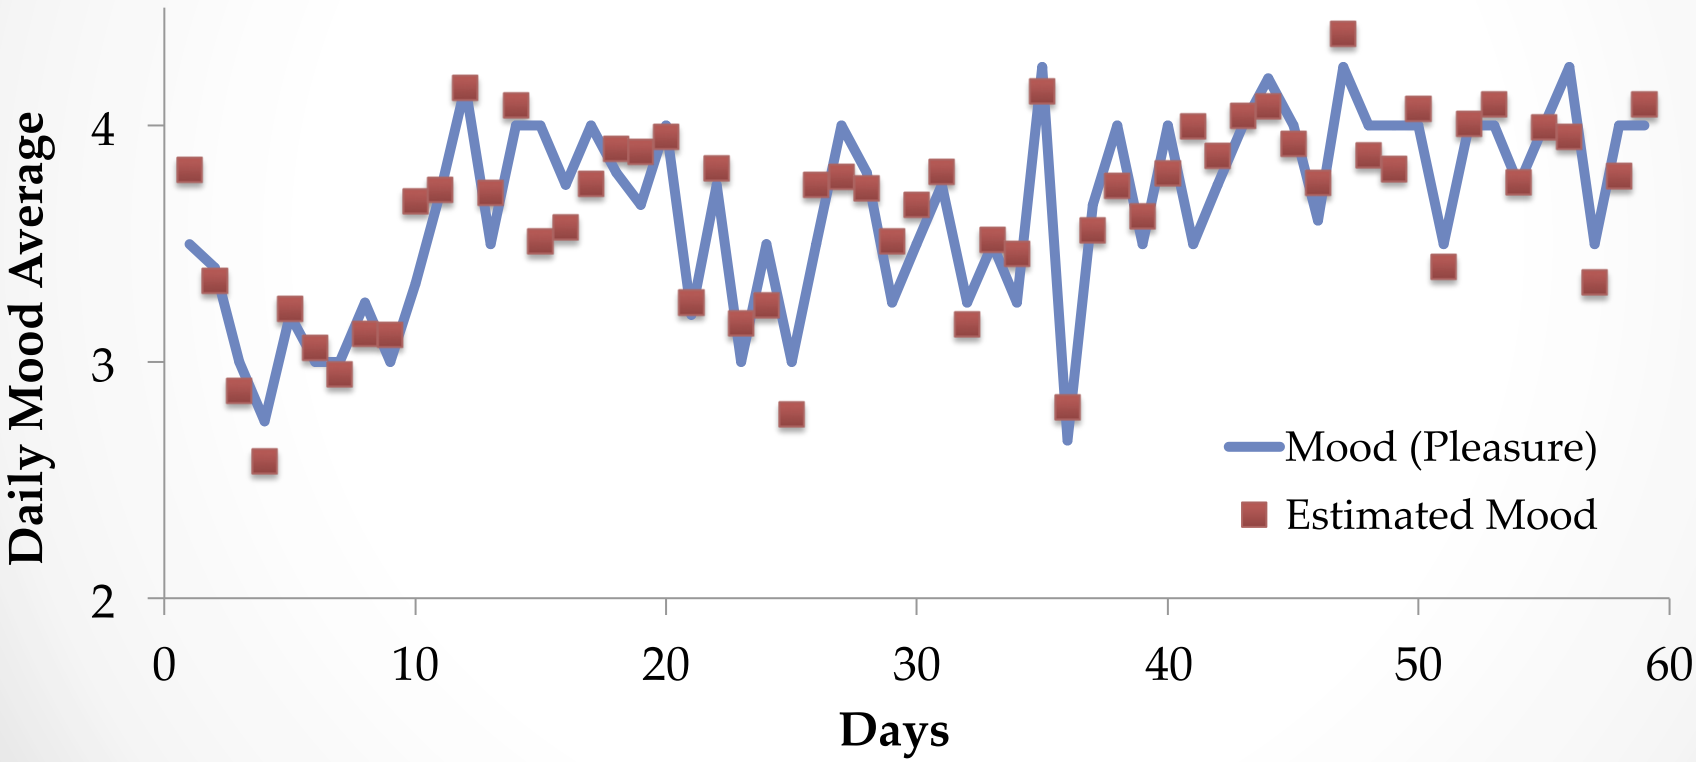

Error distributions

- Error^2 of > 0.25 will misclassify a mood label

93% < 0.25 error^2

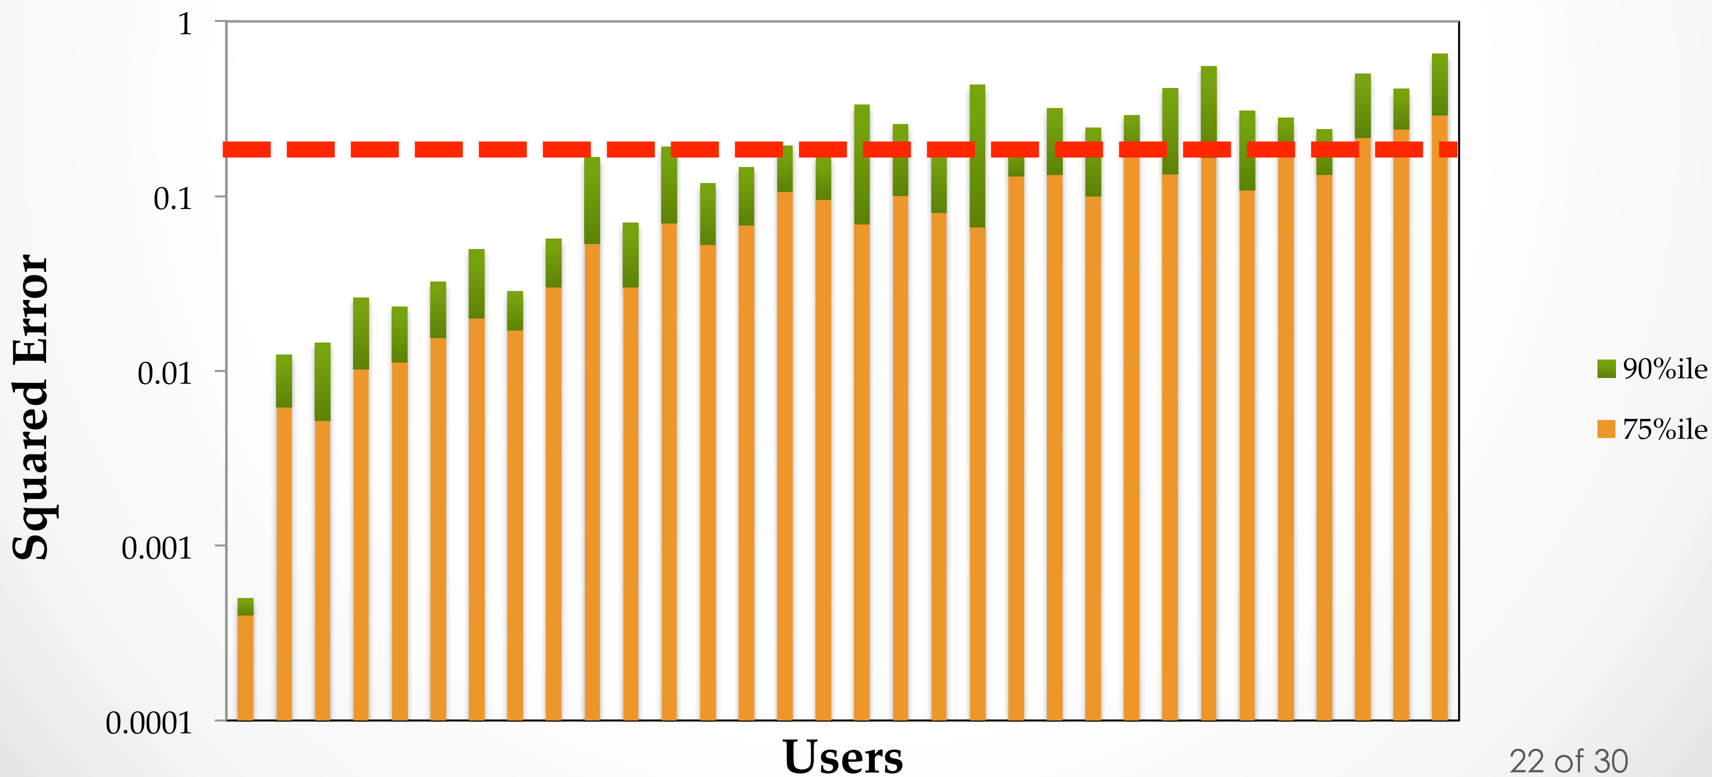

vs. Strawman Models

Models using **full-knowledge** of a user's data with LOOCV

Model A: *Assume User's Average Mood*

73% Accuracy

Model B: *Assume User's Previous Mood*

61% Accuracy

MoodScope Training: 93% Accuracy.

Personalized Training

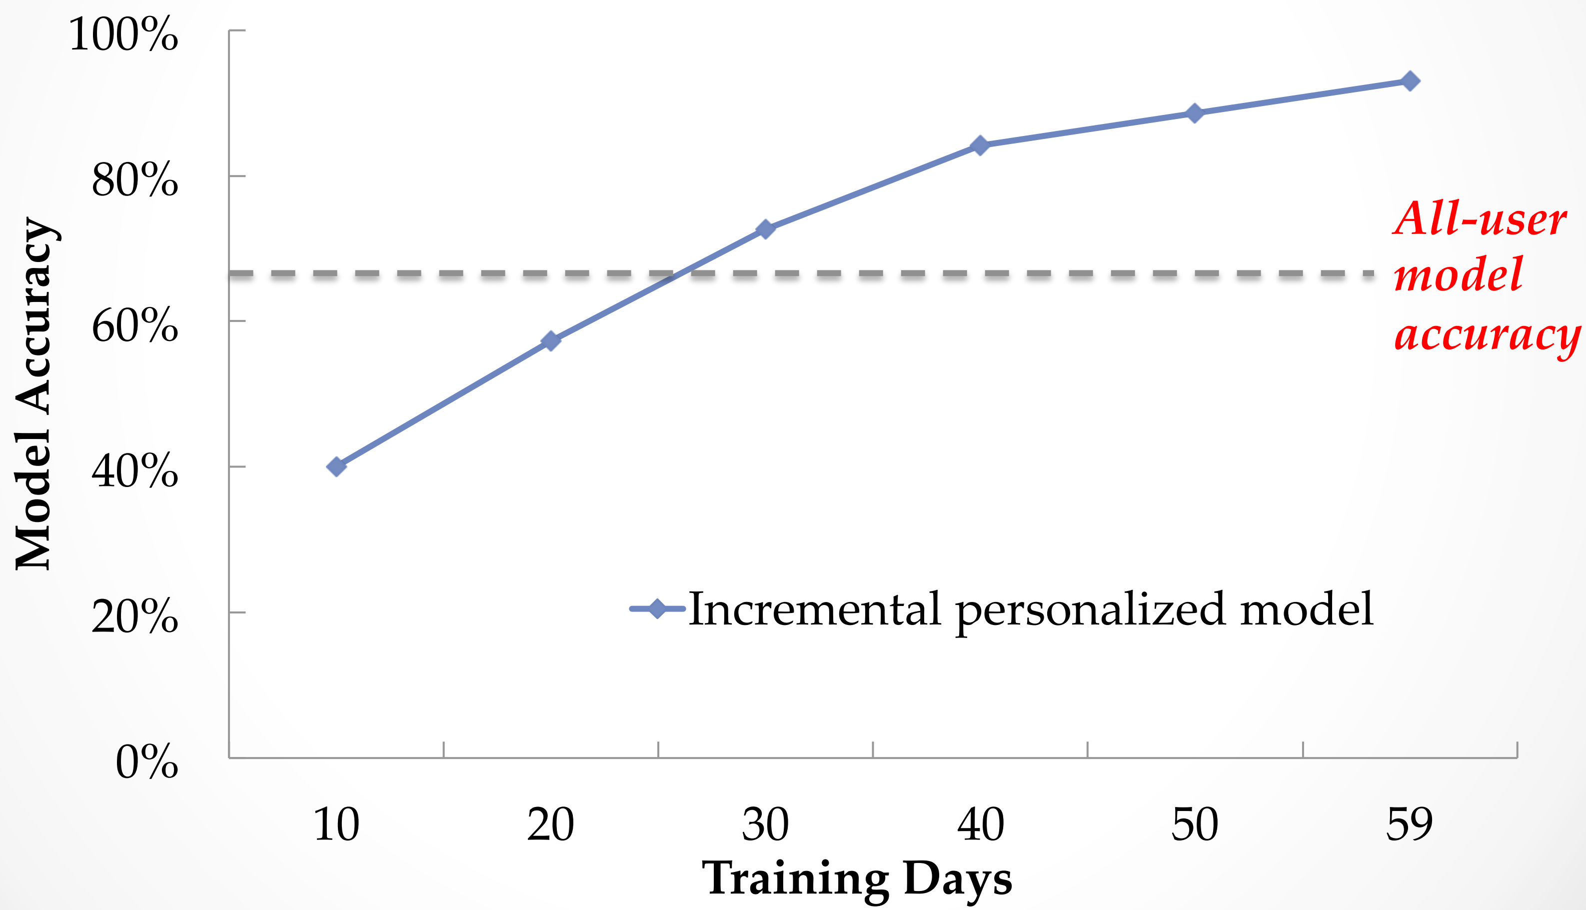

Personalized/All-user Hybrid Training

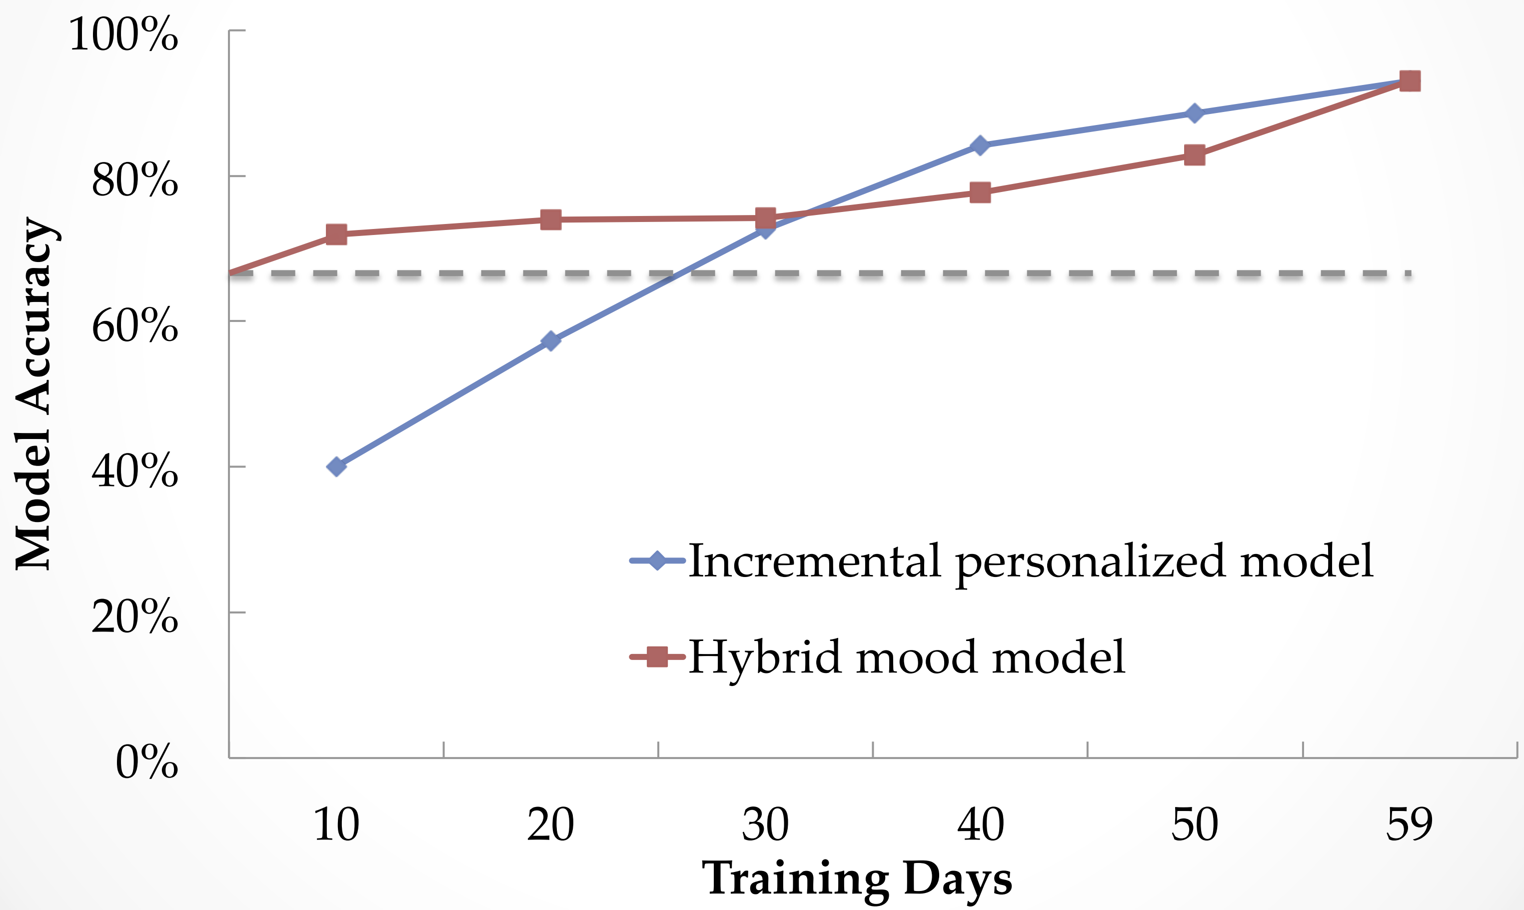

Resource-friendly Implementation

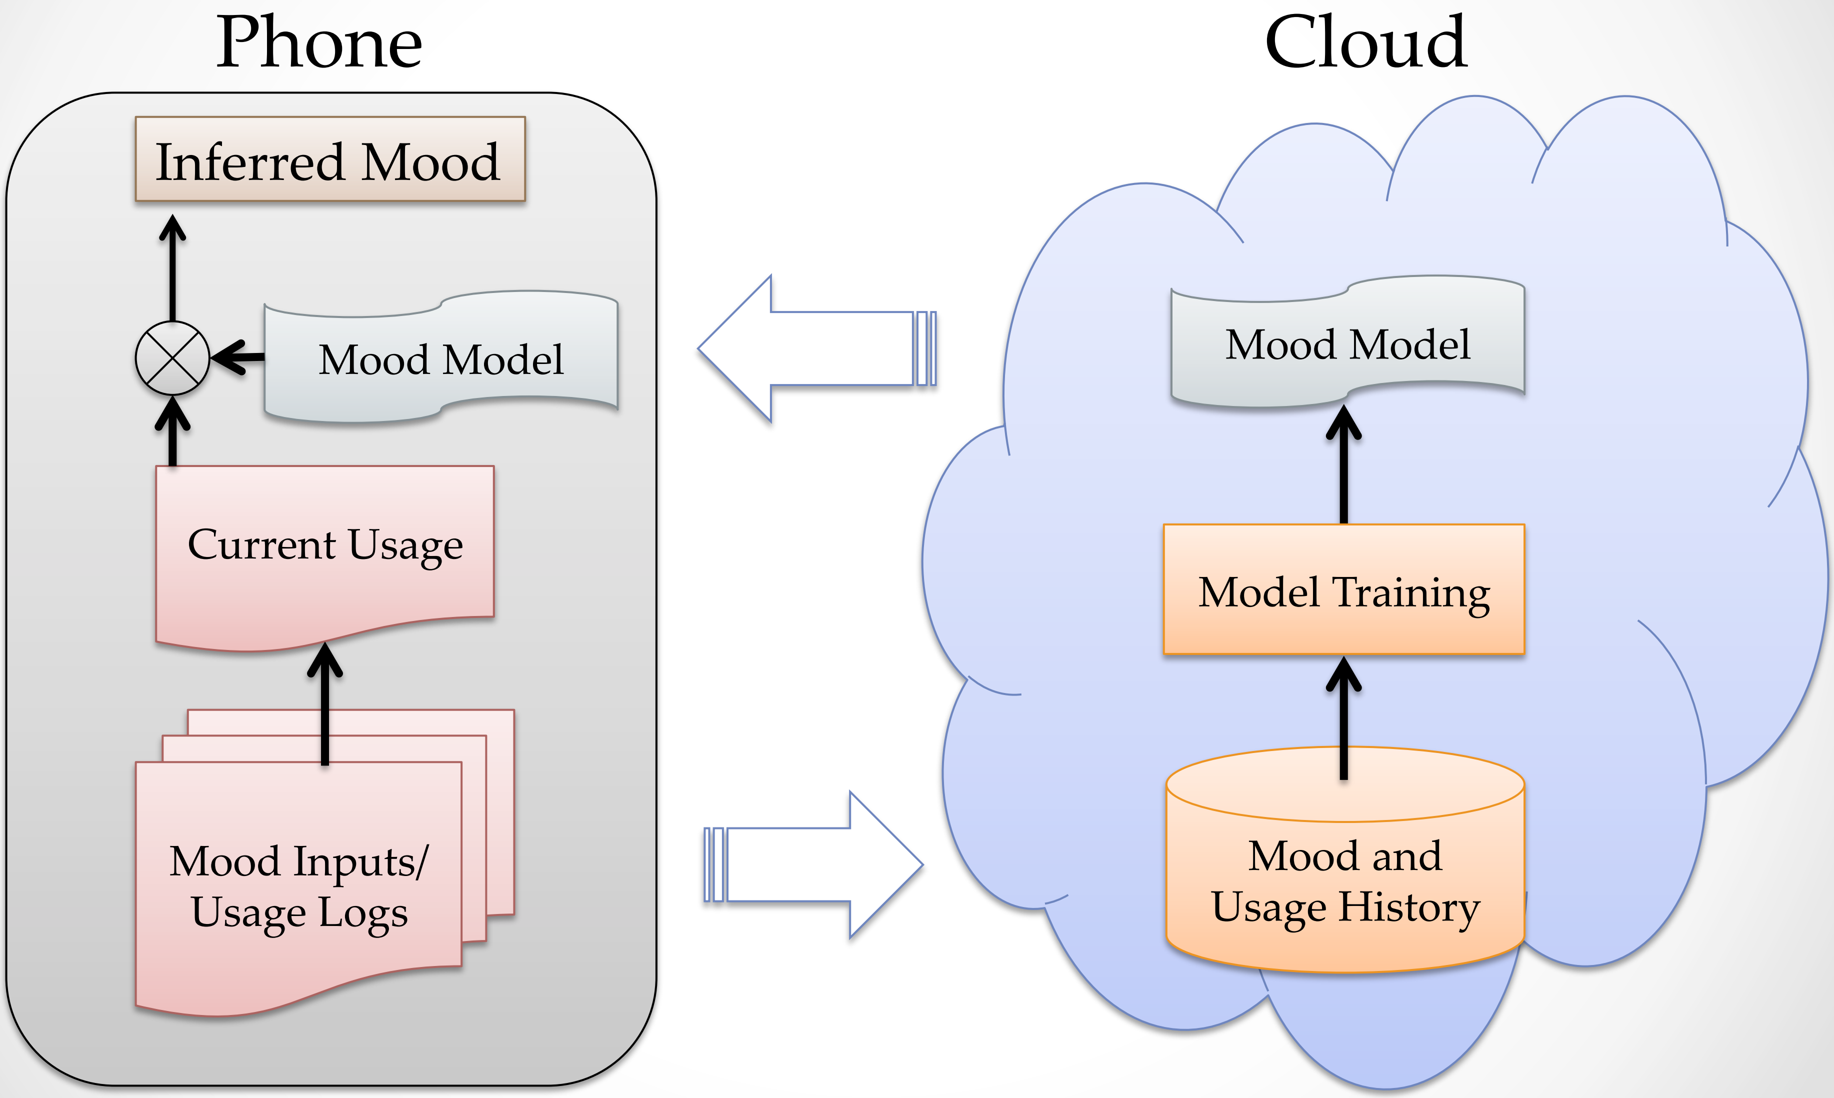

Inferred Mood

API

MoodScope:

Sensing mood from smartphone usage patterns

- Robustly (93%) detect each dimension of daily mood
    - On personalized models
    - Starts out with 66% on generalized models
- Validate with 32 users x 2 months worth of data
- Simple resource-friendly implementation