



Uncertainty Quantification for Data-Driven Decision-Making, with Applications to Weather-Sensitive Construction

Rafael Izbicki
(UFSCar)



Agradecimentos: FAPESP, CAPES, CNPq



Construction is a weather-sensitive process

Typical construction schedule: overlapping activities, task-specific weather windows, and dependencies



Shift the start date → the entire network meets a different part of the weather calendar.

Why estimate project duration after delays?

Deadline risk

Missed completion dates can trigger contractual penalties.

Start-date choice

Most construction-ML papers stop at standard regression

2009 Pewdum et al.

Forecasting final budget and duration of highway construction projects

POINT OUTPUT • duration / budget

2014 Wauters & Vanhoucke

Support vector machine regression for project control forecasting

POINT OUTPUT • regression forecast

2020 Yaseen et al.

Prediction of risk delay in construction projects using a hybrid AI model


POINT OUTPUT • predicted delay risk

2021 Egwim et al.


Applied artificial intelligence for predicting construction projects delay

POINT OUTPUT • predicted delay

Same mean, very different decisions



First, supervised learning



LLMs



Autonomous driving



AI drug discovery

Supervised Learning

Given measurements $(\mathbf{X}_1, Y_1), \dots, (\mathbf{X}_n, Y_n)$, learn a model to **predict** Y_i based on $\mathbf{X}_i$

construct g to obtain **good predictions**

$$g(\mathbf{X}_{n+1}) \approx Y_{n+1}, \dots, g(\mathbf{X}_{n+m}) \approx Y_{n+m}$$

Which one is the dog?



However, even the best models make mistakes.

How to Estimate Conditional Densities?

$$f(y|x)$$

Usual Models

$$Y \mid \mathbf{x} \sim \mathcal{N}(\beta^\top \mathbf{x}; \sigma^2)$$

$$Y \mid \mathbf{x} \sim \mathcal{N}(\mu(\mathbf{x}), \sigma(\mathbf{x}))$$

$$\hat{f}(y \mid \mathbf{x}) = \frac{\sum_{i=1}^n K_{h_x}(\|\mathbf{x} - \mathbf{X}_i\|) K_{h_y}(y - Y_i)}{\sum_{i=1}^n K_{h_x}(\|\mathbf{x} - \mathbf{X}_i\|)}$$

FlexCode

$$f(y|x) = \sum_i \beta_i(x) \phi_i(y)$$

$$\beta_i(x) = \mathbb{E}[\phi_i(Y)|x]$$

Electronic Journal of Statistics

Vol. 11 (2017) 2800–2831

ISSN: 1935-7524

DOI: [10.1214/17-EJS1302](https://doi.org/10.1214/17-EJS1302)

Converting high-dimensional regression to high-dimensional conditional density estimation^{*}

Rafael Izbicki[†]

Department of Statistics, Federal University of São Carlos, Brazil

e-mail: rafaelizbicki@gmail.com

and

Ann B. Lee[‡]

Department of Statistics, Carnegie Mellon University, USA

e-mail: annlee@cmu.edu

A First Measurement of Circumgalactic Dust Reddening from Only 4.6 deg² of the Rubin Observatory's Data Preview 1

John Franklin Crenshaw^{1,2,3} , Matthew McQuinn⁴ , and Jessica K. Werk⁴ 

¹Kavli Institute of Particle Astrophysics and Cosmology, USA; jfcren@stanford.edu

²Department of Physics, Stanford University, Stanford, CA 94305-4085, USA

³SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, CA 94025, USA

⁴Department of Astronomy, University of Washington, Seattle, WA 98195, USA

Received 2026 June 23; revised 2026 August 17; accepted 2026 September 2; published 2026 September 17

Abstract

We present the first measurement of circumgalactic dust reddening from the Vera C. Rubin Observatory, using only 4.6 deg² of ComCam Data Preview 1—roughly 0.03% of the final LSST footprint. Using photometric redshifts, we stack background-galaxy colors around foreground-galaxy positions and detect a chromatic reddening profile from $r_{\perp} \simeq 10$ kpc to 1 Mpc. Interpreting average $E(g - z)$ with a Milky Way extinction curve, we find $A_V = (1.3 \pm 0.4) \times 10^{-1} (r_{\perp}/20 \text{ kpc})^{-1.8 \pm 0.4}$ within 120 kpc. The amplitude and radial dependence agree with earlier Sloan Digital Sky Survey, KiDS, and Dark Energy Survey results despite the $\sim 1000\times$ smaller survey area and a foreground sample extending 3–6 mag fainter and 1–2 dex lower in stellar mass. The innermost 10–15 kpc bin reaches $A_V \simeq 0.3$ mag, comparable to high-latitude extinction through the Milky Way disk near the solar circle; the steep power-law slope implies a dust distribution that does not simply trace the halo-gas profile. Splitting by rest-frame $g - r$ shows stronger extinction around red foreground galaxies (rest-frame $g - r > 0.5$), although the blue subsample is too noisy to establish a significant color dependence. This red sample, with median halo mass $5 \times 10^{11} M_{\odot}$, shows substantially more reddening within 50 kpc than previously measured around more massive luminous red galaxies and implies a dust-to-stellar-mass ratio of $\sim 2\%$, nearly saturating the dust budget allowed by stellar metal yields. These pathfinder data demonstrate LSST's promise for high-precision galaxy–dust measurements across galaxy mass, environment, and redshift.



Tabular Foundation Models

nature[View](#)

[Explore content](#) ▾ [About the journal](#) ▾ [Publish with us](#) ▾

[nature](#) > [articles](#) > article

Article | [Open access](#) | Published: 08 January 2025

Accurate predictions on small data with a tabular foundation model

[Noah Hollmann](#) ✉, [Samuel Müller](#) ✉, [Lennart Purucker](#), [Arjun Krishnakumar](#), [Max Körfer](#), [Shi Bin Hoo](#), [Robin Tibor Schirrmeister](#) & [Frank Hutter](#) ✉

[Nature](#) **637**, 319–326 (2025) | [Cite this article](#)

[Save article](#)

550k Accesses | 962 Citations | 524 Altmetric | [Metrics](#)

Photometric Redshift Prediction in S-PLUS


DRAFT VERSION AUGUST 12, 2026

Typeset using L^AT_EX twocolumn style in AASTeX7.0.1

Tabu

R

⁵ Eng



US

il

Conditional Densities vs Quantile Regression

CONDITIONAL DENSITY

Full predictive distribution

QUANTILE REGRESSION

Selected quantiles only

How to Estimate Conditional Densities on the construction problem?

Modeling task workability probabilities by location and day

UK MIDAS • 1985-2020

68 stations retained

Station network



schematic spatial coverage

Task-specific no-work rules

Condition	E	C	F	S	O	P
$T_{min} \leq 0^{\circ}\text{C}$		×				×
$T_{mean} \leq 0^{\circ}\text{C}$	×			×	×	
$T_{max} \geq 40^{\circ}\text{C}$		×		×		×
$\text{Rain} \geq 1 \text{ mm}$					×	×
$\text{Rain} \geq 10 \text{ mm}$	×	×			×	×
$\text{Rain} \geq 30 \text{ mm}$	×	×		×	×	×
$\text{Gust} \geq 30 \text{ kt}$	×	×	×	×	×	

E earthworks C concrete F formwork S steelwork O outdoor painting P pavements

Weather thresholds → $H^t_{l,v,i} \in \{0,1\}$: workable / not workable

Estimating workable-day probabilities:

Binary history

1

$$H_{\ell, y, i}^{(t)} \in \{0, 1\}$$

Was task t workable at location ℓ , year y , day-of-year i ?



Smooth across space, season and years

2

$$\hat{\theta}_i^{(t)} = \frac{\sum_{\ell, i', y} w_{\ell, i', y} H_{\ell, y, i'}^{(t)}}{\sum_{\ell, i', y} w_{\ell, i', y}}$$

$$w_{\text{loc}} \propto 1/d(\ell_0, \ell)$$

$$w_{\text{year}} = \alpha^{y_0 - y}$$



$$w_{\ell, i', y} = w_{\text{loc}} w_{\text{doy}} w_{\text{year}}$$

w_{doy} : moving window around day i

Typical construction schedule: overlapping activities, task-specific weather windows, and dependencies



Shift the start date → the entire network meets a different part of the weather calendar.

■ scheduled activity □ weather requirement → dependency

Estimating workable-day probabilities:

1

Binary history

$$H_{\ell, y, i}^{(t)} \in \{0, 1\}$$

Was task t workable at location ℓ , year y , day-of-year i ?



Metrics for probabilistic forecasts

DENSITY LOSS

$$L_{\text{dens}}(\hat{f}; Y, x) = \int \hat{f}(y | x)^2 dy - 2\hat{f}(Y | x)$$


PINBALL LOSS

$$\ell_{\alpha}(y, q_{\alpha}(x)) = (\alpha - \mathbf{1}\{y < q_{\alpha}(x)\}) \times (y - q_{\alpha}(x))$$


RIGHT-TAIL PINBALL

$$L_w(Y, x) = \int_0^1 w(\alpha) \ell_{\alpha}(Y, q_{\alpha}(x)) d\alpha, \\ w(\alpha) \propto \alpha^2$$


Weather risk is task-specific and seasonal




From Uncertainty to Decisions



(c) Density



How to give a single point prediction?



Missing deadline

CRC for construction decisions

**OUTPUT: RISK-CONTROLLED CONTRACTUAL
DURATION**

$$\hat{d} \in \arg \min_{d \in \hat{\mathcal{D}}} \mathbb{E}[d(X)] \quad \text{subject to} \quad \mathbb{E}[c\{Y, d(X)\}] \leq \alpha$$

$$\hat{\mathcal{D}} = \{\hat{Q}_\lambda : \lambda \in \Lambda\}.$$


OVERRUN PROBABILITY

$$c(Y, d) = \mathbf{1}\{Y > d\}$$


OVERRUN SEVERITY

$$c(Y, d) = g((Y - d)_+), \quad 0 \leq g \leq 1, \quad g \uparrow$$


Rectified Conformal Risk Control



What can break?



2018 was a different UK weather regime




Source: Met Office summer climate summaries; anomalies relative to 1981-2010.

Take-home message

- Predict distributions when decisions depend on more than the mean
- Separate sharpness from reliability: uncertainty needs calibration
- Propagate uncertainty to the downstream decision


**Machine Learning
Beyond Point Predictions:
Uncertainty Quantification**



Rafael Izbicki

rafaelizbicki.com/UQ4ML

rafaelizbicki.com/ame/





Questions?

rafaelizbicki@gmail.com