Probabilistic Error Guarantees for Abductive Inference

Kerria Pang-Naylor   Ian Li   Kishore Rajesh
George D. Montañez

AMISTAD Lab
Department of Computer Science
Harvey Mudd College
Claremont, CA

Table of Contents

What is Abduction?

Abduction: given effects, determine most probable cause

- "Backwards" reasoning
- Unavoidable uncertainty

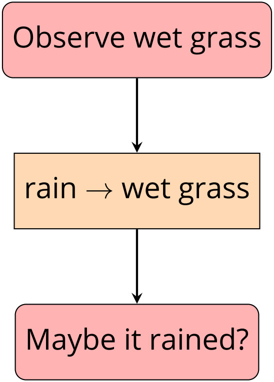

We want to *formally bound* this uncertainty.

Motivation – Abduction in Data Science

Abductive processes are ubiquitous in everyday data science techniques.

- Classifier inference
- Bayesian networks
- *maximum a posteriori* (MAP) estimation

Abductive error bounds describe these processes.

Motivation – Abduction and Intelligence

Current artificial intelligence is generally correlation-based.

- Humans reason through causal (abductive) understanding

Figure: DALL-E and MidJourney generated hands (Ragan, 2023).

Greater *quantitative* understanding of abduction needed to bridge this gap.

Overview

We present a novel framework of abduction with which we establish two sets of error bounds:

1. Assumes selection of the *single* most likely cause
   - Bayesian probability
   - Accounts for uncertainty in background information
2. Assumes selection of *any* cause whose probability is above threshold q_{min}
   - Accounts for noisy observations

Table of Contents

Setup: Posterior Probabilities

Given observation $\mathbf{x}$, what is the probability that C is the cause?

Bayes' Theorem (Posterior Probability)

$$\Pr(C|\mathbf{x}) = \frac{\Pr(\mathbf{x}|C)\Pr(C)}{\Pr(\mathbf{x})} \propto \Pr(\mathbf{x}|C)\Pr(C)$$

- **Likelihood** ($\Pr(\mathbf{x}|C)$): assuming C, probability of $\mathbf{x}$
  - e.g., $\Pr(\text{wet grass}|\text{rain})$ is very high
- **Prior** ($\Pr(C)$): overall probability of C
  - e.g., $\Pr(\text{rain})$ is low in Southern California

Any observation is represented as a binary vector whose components indicate observed events.

- **“Observation Space”($\mathcal{O}$)** of possible observations
- 2^k possible observation vectors for k events

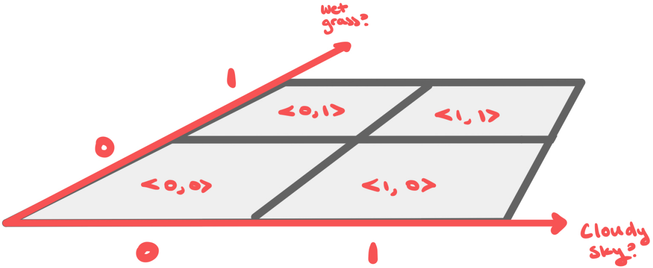

Every possible cause holds certain *likelihood* probabilities for each observation vector in O.

- 1-1 correspondence with a likelihood PMF over $\mathcal{O}$
- Posterior distribution is scaled by prior

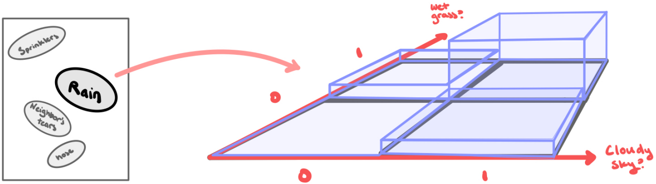

Assume *q*-percent confidence bounds for priors, likelihoods, and (thus) posteriors.

Table of Contents

Results: Bounds for Maximum Posterior Selection

Definition

Let $\mathcal{M}$ denote the set of posteriors for all causes. Each $M_i \in \mathcal{M}$ has an estimated posterior distribution with a q-probability confidence interval (l_i, u_i). Then, we can set

$$l = \max(\{l_i | i \in \mathbb{Z}, i \leq |\mathcal{M}|\})$$

and

$$u = \max(\{u_i | i \in \mathbb{Z}, i \leq |\mathcal{M}|\})$$

Selection of the Highest Posterior

Select the cause C_i whose posterior M_i is the highest of all $\mathcal{M}$.

Theorem

The probability of $M_i \in \mathcal{M}$ being the max posterior can be found as the following

$$\Pr(IsMax(M_i)) = \int_l^u \Pr\left(\bigcap_{j=1, j\neq i}^{|\mathcal{M}|} (M_j < x) \,\middle|\, M_i = x \right) p_{M_i}(x) dx$$

Bayes Error Rate

Two ways to fail when selecting the maximum posterior:

1. Incomplete or imprecise background information ($\rightarrow$ posterior *q*-percent confidence intervals)
2. The true cause is not the cause with the highest posterior ($\rightarrow$ **Bayes Error Rate**)

The error rate of selecting M_i assuming all posteriors lie in their *q*-percent confidence intervals is $\gamma_{i,upper}$ and $\gamma_{i,lower}$.

- combines Bayes Error Rates with $\Pr(\text{IsMax}(M_i))$

Uncertainty of Maximum Posterior Selection

Theorem

Let q^k be the probability that all $M_i \in \mathcal{M}$ lie in their respective confidence bounds $[l_i, u_i] \in \mathcal{U}$. The probability of wrong abduction $\Pr(W)$ *is bounded as follows:*

$$\gamma_{i,\,lower} \cdot q^k \leq \Pr(W) \leq 1 - q^k \cdot (1 - \gamma_{i,\,upper})$$

Table of Contents

Milder, Noise-accounting Bounds

Redefine success as the selection of *any* cause with estimated posterior above q_{min}.

- Reframe as **search problem** with the Algorithmic Search Framework (ASF)
- ASF property: bounds account for noisy observations

Algorithmic Search Framework (ASF)

(Montañez, 2017)

- Ω – search space (all causes)
- $T \subseteq \Omega$ – target set (causes with posteriors above q_{min})
- F – information resource (observation + estimated posterior)
- $\mathcal{A}$ – Search Algorithm (maximize posterior)

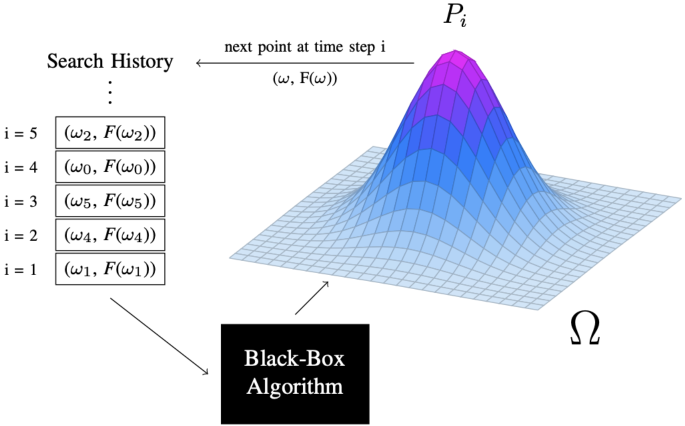

Theorem

The probability of a successful abduction, θ, is bounded above by

$$\theta \leq \frac{I(T;F) + D(P_T \| \mathcal{U}_T) + 1}{I_\Omega}.$$

Theorem

The probability of a successful abduction, θ, is bounded above by

$$\theta \leq \frac{\mathbf{I}(\mathbf{T};\mathbf{F}) + D(P_T||\mathcal{U}_T) + 1}{I_\Omega}.$$

- $I(T;F)$ – dependence between target set and observation

Bounds for Threshold Posterior Selection

Theorem of Success Under Dependence

Theorem

The probability of a successful abduction, θ, is bounded above by

$$\theta \leq \frac{I(T;F) + \mathbf{D}(\mathbf{P_T} || \mathbf{U_T}) + 1}{I_\Omega}.$$

- $I(T;F)$ – dependence between target set and observation
- $D(P_T || \mathcal{U}_T)$ – non-uniformness of target

Theorem

The probability of a successful abduction, θ, is bounded above by

$$\theta \leq \frac{I(T;F) + D(P_T||\mathcal{U}_T) + 1}{\boldsymbol{I}_\Omega}.$$

- $I(T;F)$ – dependence between target set and observation
- $D(P_T||\mathcal{U}_T)$ – non-uniformness of target
- $I_\Omega = -\log(\frac{|T|}{\Omega})$ – baseline problem difficulty

Consequence: High Likelihood Causes are Rare

Theorem

For any fixed search problem (Ω, T, F)*, set of probability mass functions* $\mathcal{P} = \{P : P \in [0,1]^{|\Omega|}, \sum_j P_j = 1\}$*, and a fixed threshold* $q_{\min} \in [0,1]$*,*

$$\frac{\mu(\mathcal{G}_{t,q_{\min}})}{\mu(\mathcal{G}_{\mathcal{P}})} \leq \frac{p}{q_{\min}},$$

where $p = \frac{|T|}{|\Omega|}$*,* $\mathcal{G}_{t,q_{\min}} = \{P : P \in \mathcal{P}, t^\top P \geq q_{\min}\}$*, and* μ *is Lebesgue measure.*

- $\frac{\mu(\mathcal{G}_{t,q_{\min}})}{\mu(\mathcal{G}_{\mathcal{P}})}$ – proportion of causes with at least $q_{\min}$ likelihood
- As $q_{\min}$ grows (high likelihood), this proportion shrinks

Table of Contents

Future Work

We formalize and derive two sets of error bounds for selective abduction. Future directions include:

- More rigorous experiments with real data
- Continuous-valued observations
- Creative abduction: *generate* cause/hypothesis given observations

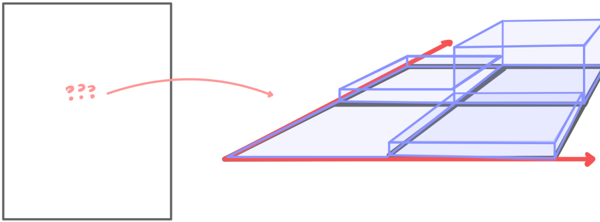

Erik J. Larson (2021)
The Myth of Artificial Intelligence: *Why Computers Can't Think the Way We Do*
The Belknap Press of Harvard University Press, 2021

George Montañez (2017)
The Famine of Forte: Few Search Problems Greatly Favor Your Algorithm
IEEE SMC 2017, 477 – 482.

Sekeh, S. Y., Oselio, B., and Hero, A. O. (2018)
Learning to Bound the Multi-class Bayes Error
IEEE Transactions on Signal Processing, 68:3793–3807

Prashnu Verma (2023)
Bad News for Deep-Fakes
Washington Post, washingtonpost.com

Acknowledgements

This work is supported by NSF grant no. 2243941 and the Harvey Mudd College Experiential Learning Fund.