

A Coupon-Collector Model of Machine-Aided Discovery*

Aditya Vempaty
IBM T. J. Watson Research Center
1101 Kitchawan Road
Yorktown Heights, New York
avempat@us.ibm.com

Lav R. Varshney
Department of Electrical and
Computer Engineering,
University of Illinois at
Urbana-Champaign,
Urbana, Illinois
varshney@illinois.edu

Pramod K. Varshney
Department of Electrical Engineering
and Computer Science,
Syracuse University,
Syracuse, New York
varshney@syr.edu

ABSTRACT

Empirical studies of scientific discovery—so-called Eurekometrics—have indicated that the output of exploration proceeds as a logistic growth curve. Although logistic functions are prevalent in explaining population growth that is resource-limited to a given carrying capacity, their derivation do not apply to discovery processes. This paper develops a generative model for logistic *knowledge discovery* using a novel extension of coupon collection, where an explorer interested in discovering all unknown elements of a set is supported by technology that can respond to queries. This discovery process is parameterized by the novelty and quality of the set of discovered elements at every time step, and randomness is demonstrated to improve performance. Simulation results provide further intuition on the discovery process.

KEYWORDS

probabilistic modeling, generalized coupon collector’s problem, knowledge discovery

1 INTRODUCTION

Eurekometric studies quantify the process of discovery and have found that the empirical number of scientific discoveries as a function of time initially increases exponentially before saturating to a *capacity* limit according to an approximate logistic curve [2–4]. This is seen in the discovery of mammalian species, chemical elements, and minor planets [2]. The logistic curve is often used to describe population growth that is resource-limited to a given carrying capacity [16], but population growth dynamics do not seem to describe scientific discovery. This paper develops a generative model for discovery that recreates empirical observations, by generalizing the coupon collector model [5, 7, 10, 14]. In particular, the units of discovery (species, elements, planets) are the coupons and an explorer is trying to collect (discover) them all.

There are two main generalizations introduced into the standard coupon collector formulation.

Scientific discovery problems can be cast as big-D discovery,¹ where there is also growing interest in developing artificial intelligence (AI) support [9, 11, 18]. One can, however, just as easily think about little-d knowledge discovery problems such as looking for

restaurants on Yelp,² where one might ‘discover’ a new restaurant. In addition to being novel, the restaurant should also be of high quality and have good ratings. In the sequel, we will therefore evaluate the discovery problem using notions of both novelty and quality, e.g. as judged by a suitably knowledgeable social group [15].

In the economics, cognitive psychology, and AI literatures, problem solving is described in terms of searching a problem space, where the first step is discovery to determine available elements in the space [1]. In rationality-based economic theory, the set of possible alternatives is known [13], but under the bounded rationality framework [13, 17], the set of alternatives is not known *a priori*. Thus limited agents must perform discovery, perhaps with technological support to overcome limitations.³ In scientific discovery, new technologies for separating compounds into constituent parts enable discovery of new chemical elements and new telescopes enable discovery of new minor planets, for example. In the sequel, we will consider discovery as supported by a technology, which may itself be noisy.

2 PROBLEM FORMULATION

The discovery problem is modeled as follows. Let the universe of M total elements be denoted Θ . The explorer is initially only aware of a subset $\Theta_0 \subset \Theta$ of these M possibilities. She learns of other elements in steps, where at each step she learns of a new object. In taking the step, a previously unknown object is drawn from the M possible values according to the prior probability mass function (pmf) $\mathbf{p} = [p_1, \dots, p_M]$, but the supporting technology yields noisy observations of which kind of object it is. The explorer updates her element set accordingly. Note that the technology presents its outcome to the explorer without any information regarding the initial known subset of elements Θ_0 . In particular, let the underlying element θ_t be investigated at time step t . The technology makes an observation x_t^M corrupted by noise, which is i.i.d. across time, and infers the element to be $\hat{\theta}_t \in \Theta$ as follows:

$$\hat{\theta}_t = \arg \max_{\theta \in \Theta} p(\theta | x_t^M),$$

where $p(\theta | x_t^M)$ is the posterior distribution given the observation. The explorer’s updated element set is $\Theta_t = \Theta_{t-1} \cup \hat{\theta}_t$. Note that a new element gets added only if it is not already known by the explorer, who can access the true value. Fig. 1 summarizes this technology-aided discovery process.

*This work was supported in part by the IBM-Illinois Center for Cognitive Computing Systems Research (C3SR), a research collaboration as part of the IBM Cognitive Horizons Network.

¹This terminology is inspired by the terms *little-c creativity* which means novelty with respect only to the mind of the individual concerned and *big-C Creativity* which means novelty with respect to the whole of previous history or eminent creativity [6].

²<http://www.yelp.com/>

³In computational creativity, a technology supporting people in determining the set of alternatives is called a *coach* [12].




Figure 1: Single iteration of the discovery process, where $\Theta_{t-1} \subset \Theta$ is known at t , technology provides a noisy estimate $\hat{\theta}$ of unknown element θ , and $\hat{\theta}$ is discovered by the explorer if not already known.

Each new element has two basic characteristics: novelty and quality. These are detailed in turn in the next sections, in terms of the size and total quality of the discovered set.

3 SIZE OF KNOWLEDGE BASE

Size is measured by the total number of distinct elements that are known to the explorer at any given time. This includes both the prior set of elements known before the start of the discovery process and the elements discovered during discovery. Let $N_t = |\Theta_t|$. Clearly the learning rate depends on the pmf \mathbf{p} , the initial set Θ_0 , and the noise distribution. While performance is best for uniform prior when the initial knowledge set Θ_0 is arbitrary, for cases when Θ_0 contains rare elements, knowledge discovery is fast since the unknown elements are of high probability and can be quickly found by the explorer. This intuition is also observed in our analytical results. Also, for noisy observations, it is possible that higher variance yields faster total discovery as noise can cause false decisions in the technology aiding the discovery of new elements (typically of low probability) by the explorer.

Let the misclassification probabilities due to noise-corrupted observations at any time t be $r_{mn} \triangleq \Pr(\hat{\theta}_t = n | \theta_t = m)$ for $m, n = 1, \dots, M$. Therefore, the posterior distribution $\tilde{\mathbf{p}} = [\tilde{p}_1, \dots, \tilde{p}_M]$ of the elements is the new prior distribution of the explorer's input. An asymmetric channel $\mathbf{R} \triangleq \{r_{mn}\}$ represents the difficulty of identifying elements. As has been noted in empirical studies, some elements may be more difficult to identify/discover than others [2–4].

PROPOSITION 3.1. *For the technology-aided discovery problem with a pmf \mathbf{p} and noise channel \mathbf{R} , the expected number of elements learned after T time steps is*

$$E[N_T | \Theta_0] = |\Theta_0| - \sum_{k=1}^T \sum_{\theta \notin \Theta_0} (-1)^k \binom{T}{k} \tilde{p}_\theta^k. \quad (1)$$

PROOF. The number of elements are updated as follows

$$N_t = \begin{cases} N_{t-1} & \text{if } \hat{\theta}_t \in \Theta_{t-1} \\ N_{t-1} + 1 & \text{if } \hat{\theta}_t \notin \Theta_{t-1}, \end{cases} \quad (2)$$

resulting in a Markov Chain:

$$N_t = \begin{cases} N_{t-1} & \text{with probability } \sum_{\theta \in \Theta_{t-1}} \tilde{p}_\theta \\ N_{t-1} + 1 & \text{with probability } 1 - \sum_{\theta \in \Theta_{t-1}} \tilde{p}_\theta. \end{cases} \quad (3)$$

Therefore,

$$\begin{aligned} E[N_t | \Theta_{t-1}] &= |\Theta_{t-1}| \sum_{\theta \in \Theta_{t-1}} \tilde{p}_\theta + (|\Theta_{t-1}| + 1)(1 - \tilde{p}_\theta) \\ &= |\Theta_{t-1}| + 1 - \sum_{\theta \in \Theta_{t-1}} \tilde{p}_\theta. \end{aligned} \quad (4)$$

For $t = 1$,

$$E[N_1 | \Theta_0] = m_0 + 1 - \sum_{\theta \in \Theta_0} \tilde{p}_\theta, \quad (5)$$

and for $t = 2$,

$$E[N_2 | \Theta_1] = N_1 + 1 - \sum_{\theta \in \Theta_1} \tilde{p}_\theta = N_1 + 1 - \sum_{\theta \in \Theta_0} \tilde{p}_\theta - \sum_{\theta \in \Theta_1 \setminus \Theta_0} \tilde{p}_\theta. \quad (6)$$

Recursively, we have for $t \geq 2$,

$$\begin{aligned} E[N_t | \Theta_{t-2}] &= E \left[N_{t-1} + 1 - \sum_{\theta \in \Theta_{t-2}} \tilde{p}_\theta - \sum_{\theta \in \Theta_{t-1} \setminus \Theta_{t-2}} \tilde{p}_\theta \middle| \Theta_{t-2} \right] \\ &= |\Theta_{t-2}| + 2 \sum_{\theta \notin \Theta_{t-2}} \tilde{p}_\theta - E \left[\sum_{\theta \in \Theta_{t-1} \setminus \Theta_{t-2}} \tilde{p}_\theta \middle| \Theta_{t-2} \right]. \end{aligned} \quad (7)$$

An additional element is added into Θ_t only when it is not originally present in Θ_{t-1} , which is with probability $\sum_{\theta \notin \Theta_{t-1}} \tilde{p}_\theta$. Under this condition that an element has been added, each $\theta \notin \Theta_{t-2}$ occurs with probability $\tilde{p}_\theta / \sum_{\theta \notin \Theta_{t-2}} \tilde{p}_\theta$, which reduces (7) to the following

$$E[N_t | \Theta_{t-2}] = |\Theta_{t-2}| + 2 \sum_{\theta \notin \Theta_{t-2}} \tilde{p}_\theta - \sum_{\theta \notin \Theta_{t-2}} \tilde{p}_\theta^2. \quad (8)$$

Continuing further, for a general T

$$E[N_T | \Theta_0] = |\Theta_0| - \sum_{k=1}^T \sum_{\theta \notin \Theta_0} (-1)^k \binom{T}{k} \tilde{p}_\theta^k. \quad (9)$$

□

Some points to be noted from Proposition 3.1 include the following:

- If rare elements are already known, the convergence is faster, cf. (1).
- For uniformly distributed prior pmf and perfect observations,

$$E[\rho_T | \rho_0] = 1 - (1 - \rho_0) \left(1 - \frac{1}{M} \right)^T, \quad (10)$$

where $\rho_t = N_t / M$.

- Rate of knowledge discovery (rate of convergence of $E[\rho_T | \rho_0]$ to 1) is monotonically increasing in ρ_0 .
- Growth rate of knowledge discovery is exponential in T :

$$\lim_{T \rightarrow \infty} \frac{1}{T} \log(1 - E[\rho_T | \rho_0]) = -\log \left(1 - \frac{1}{M} \right). \quad (11)$$

Therefore, the convergence is faster as M decreases.

To gain some further insight, we perform simulations. The following parameters are used:

$$\begin{aligned} \Theta &= \{1, \dots, M\}, \\ p_\theta(m) &= \binom{M-1}{m-1} p^{m-1} (1-p)^{M-m} \end{aligned}$$

for $m = 1, \dots, M$, and the noise channel is an M -ary symmetric channel with crossover probability r . Fig. 2 shows the numerical and simulation results that corroborates our analytical expressions. The simulation results are averaged over $N_{mc} = 500$ Monte Carlo




Figure 2: Numerical and simulations results of fraction of element set with iterations ($M = 4$, $\Theta_0 = [1, 2]$, $p = 0.2$, and $r = 0.1$).




Figure 3: Performance variation with different initial sets Θ_0 ($M = 4$, $p = 0.2$, and $r = 0.1$).

runs. As can be observed from Fig. 3, the performance depends on the initial set Θ_0 and the performance (in terms of discovery rate) is higher (and also better than the uniform case) when the initial elements are the least probable ones ($\Theta_0 = [3, 4]$ in this example). This is intuitive since if the explorer already knows the rare elements, she can discover others at a faster rate.

Fig. 4 shows the performance with varying noise value r where higher r implies noisier data produced by the technology. When the explorer is already aware of the most probable elements (refer to Fig. 4), noise has the positive effect of helping the explorer in discovering the less probable elements at a faster rate, implying that *noisy technology can help in discovering new elements*. Such an observation is conceptually related to the phenomenon of noise-enhanced signal processing [8], where the addition of noise can improve the performance of some non-linear suboptimal inference systems. Further analysis is needed to mathematize this relation and understand the *optimal* noise to yield best performance.

Fig. 6 shows how the performance varies with prior distribution. The effect of prior distribution on the performance is characterized by (1), but has been omitted for brevity.

For the simple case, we evaluate the performance with varying M and ρ_0 in Fig. 7 and Fig. 8, respectively.




Figure 4: Effect of noise ($M = 4$, and $p = 0.2$).




Figure 5: Effect of noise ($M = 4$, $p = 0.2$, and $\Theta_0 = [1, 2]$).




Figure 6: Effect of different prior distribution ($M = 4$, $r = 0$, and $\Theta_0 = [3, 4]$).

4 QUALITY OF KNOWLEDGE BASE

In many settings, not only is the discovery of a new element important but so is the quality of that element. For example, certain new metals are more useful in metallurgy than others. Let each element $\theta \in \Theta$ have a corresponding quality factor q_θ which determines the value of discovering θ . Let $\mathbf{q} = [q_1, \dots, q_M]$ be the quality vector and $Q_t = \sum_{\theta \in \Theta_t} q_\theta$ denote the quality of the elements discovered after time t .

Definition 4.1. The k th-order quality-prevalence function $D_{\Theta_0}^k(\tilde{\mathbf{p}}, \mathbf{q})$ is defined as the inner product between \mathbf{q} (quality vector) and




Figure 7: Expected fraction of element set with iterations ($\rho_0 = 0.5$, varying M)




Figure 8: Expected fraction of element set with iterations ($M = 10$, varying ρ_0)

element-wise k th power of $\tilde{\mathbf{p}}$ (probability vector) over the subset Θ_0^C ($\Theta \setminus \Theta_0$):

$$D_{\Theta_0}^k(\tilde{\mathbf{p}}, \mathbf{q}) := \sum_{\theta \notin \Theta_0} q_\theta \tilde{p}_\theta^k.$$

These functions evaluate the degree of alignment between the probability vector and the quality vector.

PROPOSITION 4.2. *For technology-aided discovery, the expected quality of elements discovered after T iterations is given by*

$$E[Q_T | \Theta_0] = Q_0 - \sum_{k=1}^T (-1)^k \binom{T}{k} D_{\Theta_0}^k(\tilde{\mathbf{p}}, \mathbf{q}). \quad (12)$$

PROOF. The proof is similar to that of Prop. 3.1 with the following weighted form

$$\begin{aligned} E[Q_t | \Theta_{t-1}] &= Q_{t-1} \sum_{\theta \in \Theta_{t-1}} \tilde{p}_\theta + \sum_{\theta \notin \Theta_{t-1}} (Q_{t-1} + q_\theta) \tilde{p}_\theta \\ &= Q_{t-1} + \sum_{\theta \notin \Theta_{t-1}} q_\theta \tilde{p}_\theta - \sum_{\theta \in \Theta_{t-1}} q_\theta \tilde{p}_\theta. \end{aligned} \quad (13)$$

Therefore, for a general T , we get the desired result

$$E[Q_T | \Theta_0] = Q_0 - \sum_{k=1}^T \sum_{\theta \notin \Theta_0} (-1)^k \binom{T}{k} q_\theta \tilde{p}_\theta^k. \quad (14)$$

□




Figure 9: Expected quality with time ($M = 4$, $r = 0.1$, $p = 0.2$, and $\Theta_0 = [1, 2]$).

The two extreme possibilities are when \mathbf{q} is aligned in the same direction as probability vector \mathbf{p} or is in the opposite direction. For the previous example, Fig. 9 confirms our understanding in showing the quality of the discovered elements for these two extreme cases.

5 CONNECTING TO EMPIRICAL OBSERVATIONS

As mentioned before, EurekaMetrics studies the nature of discovery [2–4]. Empirical observations in this field have shown the discovery of scientific output to follow an exponential increase, or more properly, a logistic growth curve. In our generative model, the ease of discovery corresponds to the pmf \mathbf{p} that denotes the difficulty associated with discovering an element. When all elements are of equal difficulty, empirical results suggest a logistic curve for the number of discovered elements D_T [2]:

$$D_T \approx \frac{K}{1 + Ae^{-r_0 T}} \quad (15)$$

where K is the limiting size, A is the fitting constant, and r_0 is the growth rate of scientific output. For a small value of A , this can be approximated as

$$D_T \approx K(1 - Ae^{-r_0 T}) \iff \rho_T^{emp} \approx 1 - Ae^{-r_0 T} \quad (16)$$

where $\rho_T^{emp} = D_T/K$ is the fraction of discovered elements. Observe that (16) matches (10) derived for the generative model of the discovery process.

6 DISCUSSION

The discovery problem considered herein is a novel generalization of the coupon collector’s problem [5, 7, 10, 14]. Most prior results for weighted coupon collector’s problem (or coupon collector problem in general) consider the expected number of iterations required to collect all coupons, whereas here we evaluate the average number of coupons collected after T iterations. Also, most of the existing results of coupon collectors problem are approximations/asymptotic order results. Moreover, to the best of our knowledge, there are no previous results on the coupon collector’s problem in the presence of noisy observations or where each coupon is of different quality.

The generative model for the discovery problem has also been simulated and several intuitive observations have been made. The

coupon collector's model helps interpret several empirical observations made in Eurekometrics [3] and provides further insights into the discovery process, e.g. how discoveries become more difficult over time.

REFERENCES

- [1] John R. Anderson. 2010. *Cognitive Psychology and Its Implications*. Worth Publishers, New York.
- [2] Samuel Arbesman. 2011. Quantifying the ease of scientific discovery. *Scientometrics* 86, 2 (June 2011), 245–250. <https://doi.org/10.1371/journal.pcbi.1002072>
- [3] Samuel Arbesman and Nicholas A. Christakis. 2011. Eurekometrics: Analyzing the nature of discovery. *PLoS Comput. Biol.* 7, 6 (June 2011), 1–2. <https://doi.org/10.1371/journal.pcbi.1002072>
- [4] Samuel Arbesman and Gregory Laughlin. 2010. A scientometric prediction of the discovery of the first potentially habitable planet with a mass similar to earth. *PLoS ONE* 5, 10 (Oct. 2010), 1–4. <https://doi.org/10.1371/journal.pone.0013061>
- [5] Petra Berenbrink and Thomas Sauerwald. 2009. The weighted coupon collector's problem and applications. In *Computing and Combinatorics*, Hung Q. Ngo (Ed.). Springer, Berlin-Heidelberg, 449–458.
- [6] Margaret A. Boden. 2004. *The Creative Mind: Myths and Mechanisms*. Routledge, New York.
- [7] Arnon Boneh and Micha Hofri. 1997. The coupon-collector problem revisited- a survey of engineering problems and computational methods. *Stochastic Models* 13, 1 (Jan. 1997), 39–66. <https://doi.org/10.1080/15326349708807412>
- [8] Hao Chen, Lav R. Varshney, and Pramod K. Varshney. 2014. Noise-Enhanced Information Systems. *Proc. IEEE* 102, 10 (Oct. 2014), 1607–1621. <https://doi.org/10.1109/JPROC.2014.2341554>
- [9] Ross D. King, Jem Rowland, Wayne Aubrey, Maria Liakata, Magdalena Markham, Larisa N. Soldatova, Ken E. Whelan, Amanda Clare, Mike Young, Andrew Sparkes, Stephen G. Oliver, and Pinar Pir. 2009. The robot scientist Adam. *Computer* 42, 8 (Aug. 2009), 46–54. <https://doi.org/10.1109/MC.2009.270>
- [10] John E. Kobza, Sheldon H. Jacobson, and Diane E. Vaughan. 2007. A survey of the coupon collector's problem with random sample sizes. *Methodol. Comput. Appl. Prob.* 9, 4 (Dec. 2007), 573–584. <https://doi.org/10.1007/s11009-006-9013-3>
- [11] Pat Langley. 2000. The computational support of scientific discovery. *Int. J. Human-Computer Studies* 53, 3 (Sept. 2000), 393–410. <https://doi.org/10.1006/ijhc.2000.0396>
- [12] Todd Lubart. 2005. How can computers be partners in the creative process: Classification and commentary on the Special Issue. *Int. J. Human-Computer Studies* 63, 4–5 (Oct. 2005), 365–369. <https://doi.org/10.1016/j.ijhcs.2005.04.002>
- [13] James G. March and Herbert A. Simon. 1958. *Organizations*. Wiley, New York.
- [14] K. L. Q. Read. 1998. A lognormal approximation for the collector's problem. *The American Statistician* 52, 2 (1998), 175–180. <https://doi.org/10.1080/00031305.1998.10480558>
- [15] R. K. Sawyer. 2012. *Explaining Creativity: The Science of Human Innovation*. Oxford University Press, Oxford.
- [16] Nathan F. Sayre. 2008. The genesis, history, and limits of carrying capacity. *Annals of the Association of American Geographers* 98, 1 (2008), 120–134. <https://doi.org/10.1080/00045600701734356>
- [17] Herbert Alexander Simon. 1982. *Models of Bounded Rationality: Empirically grounded economic reason*. MIT Press, Cambridge.
- [18] Raúl E. Valdés-Pérez. 1999. Principles of human-computer collaboration for knowledge discovery in science. *Artif. Intell.* 107, 2 (Feb. 1999), 335–346. [https://doi.org/10.1016/S0004-3702\(98\)00116-7](https://doi.org/10.1016/S0004-3702(98)00116-7)