

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Thompson sampling for species discovery

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1611616> since 2016-11-12T21:15:28Z

Publisher:

Electronic

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Thompson sampling for species discovery

Thompson sampling per la scoperta di nuove specie

M. Battiston, S. Favaro and Y.W. Teh

Abstract This work proposes a new methodology for discovering new species, when observations are sampled from different populations. Using a metaphor, we imagine J populations of animals to be available and we can sequentially choose from which of these populations to collect further samples. Both labels and frequencies of these species are unknown a priori. At each time step, the proposed strategy suggests where to collect the next observation in order to maximize the number of total species observed. This strategy is based on a joint use of the Hierarchical Pitman-Yor process, to estimate the unknown distributions of animals, and of Thompson Sampling for the sequential allocation problem. Performances of the algorithm are compared to those of other three strategies through simulations.

Abstract *Questo lavoro propone una nuova metodologia per scoprire nuove specie, quando le osservazioni sono campionate da diverse popolazioni. Utilizzando una metafora, immaginiamo J popolazioni di animali essere presenti e che si possa estrarre campioni da quest'ultime in maniera sequenziale. I nomi e le frequenze di ogni specie non sono note a priori. Ad ogni passo, la strategia proposta suggerisce da quale popolazione estrarre l'osservazione successiva, in modo da massimizzare il numero totale di specie osservate. Questa strategia si basa su un'uso congiunto dello Hierarchical Pitman-Yor process, per stimare le distribuzioni di animali nelle popolazioni, e del Thompson Sampling per il problema di selezionare sequenzialmente dove allocare ulteriori risorse. Le prestazioni dell'algoritmo sono comparate con quelle di altre tre strategie attraverso simulazioni.*

Marco Battiston

Bocconi University, via Sarfatti 25, Milan. e-mail: marco.battiston@phd.unibocconi.it

Stefano Favaro

University of Turin, Corso Unione Sovietica 218bis, Turin. e-mail: stefano.favaro@unito.it

Yee Whye Teh

University of Oxford, 24-29 St Giles', Oxford. e-mail: y.w.teh@stats.ox.ac.uk

Key words: Discovery probability; Thompson Sampling; Hierarchical Pitman-Yor process; Species sampling models

1 Introduction

Species sampling problems initially arose in the ecological and biological literature long time ago, but have recently found a renewed interest, due to applications in many other fields, like machine learning, genetics and linguistics. All these problems can be described using a species metaphor. We imagine a (discrete) population G of animals to be available and a sample of size n is collected from it. Each observation represents an animal and its observed value is its own species. Given this sample, interest usually lies either in predicting the realization of future observations or in estimating some particular feature of G . The frequentist literature on these problems is old and rich and a complete review on it is [3]. Whereas, contributions to the problem using a Bayesian approach are relatively recent, see e.g. [9], [4], [5] and [6].

In this work, we consider a different setting, in which J populations are available, $\{G_1, \dots, G_J\}$, and we can sequentially choose from which of these populations to collect further samples. We assume the populations to share the same species of animals and their labels and frequencies are unknown a priori. Also, the frequency of each species may vary across populations. Hence, when sampled, some populations will be more likely to display new species than others. In this work, we propose a sequential rule that, at every time step, suggests a population from which to collect the next observation, when the goal is maximizing the total number of species discovered.

The sequential allocation problem analysed in this work resembles the so-called multi-armed bandit problem, popular in the reinforcement learning literature. In the latter problem, J unknown rewards distributions are available and, at each time step, a player must choose which of these distributions to sample next, in order to maximize the total reward obtained. Inherent to this problem is a trade-off between exploiting distributions that performed well in the past or exploring those which have not been much played. In the reinforcement learning literature, the two most popular sequential strategies are the Upper Confidence Bound, which is a frequentist algorithm initially introduced in [8] and further developed by [1], and Thompson sampling (TS), which is a Bayesian procedure proposed by [15]. TS prescribes to assign priors to the unknown parameters of the J distributions and, at every time step, to play an arm according to its posterior probability of being the best one.

The problem analysed in this work can be traced back to a similar bandit formalization. Indeed, we can think a reward of one unit to be collected any time we observe a new species and zero otherwise. Hence, rewards are Bernoulli distributed with parameters varying over time, since every time a new species is observed, the probability of observing another one in the next steps decreases. Our proposed solution to the problem is made of two elements: a Bayesian nonparametric procedure

for the estimation of the $\{G_1, \dots, G_J\}$ using a Hierarchical Pitman-Yor (HPY) prior, together with a TS strategy for the sequential choice of the best arm. We refer to this strategy as the HPY-TS algorithm.

The structure of the article is as follows. In next section, we recall the HPY model and present the HPY-TS algorithm. In Section 3, we compare the performances of the HPY-TS algorithm against three competing alternatives in different scenarios.

2 The HPY-TS

The problem of sequential species discovery in presence of many populations can be described as follows. Let $\{G_1, \dots, G_J\}$ denote J discrete distributions, supported on the same countable set, assumed to be unknown a priori. The mass functions of each of these J distribution are assumed unknown too and the mass of each support point can vary across different distributions. We assume that, at each time step, we can choose one of these distributions to sample the next observation from. As mentioned before, the goal is observing as many distinct points as possible from the unknown common support.

This problem can be traced back to a stochastic bandit problem by regarding a discovery as a unitary reward. The reward of arm j at time t is then Bernoulli distributed with parameter the missing mass (i.e. the sum over all masses of the not already observed support points) of that arm at that time point. To solve this allocation problem, we proposed a strategy based on TS. TS prescribes to assign priors to the unknown parameters of each distribution and, at each time step, to play an arm according to its posterior probability of being the best one. This can be easily implemented by computing at each time the posteriors of the J mean rewards, drawing a sample from each of them and then play the arm corresponding to the highest realisation.

In our setting, the unknown parameters are the whole distributions $\{G_1, \dots, G_J\}$ and we assign a HPY prior to them. This prior for $\{G_1, \dots, G_J\}$ clearly induces a prior also for the J missing masses. Given this prior choice and given a set of data from $\{G_1, \dots, G_J\}$, we derive the joint posterior of the J missing masses, which is used to implement TS. We derive this posterior distribution and we describe the algorithm in subsection, after briefly reviewing the HPY model. We refer to proposed strategy as HPY-TS algorithm.

2.1 The HPY model

The HPY is a model introduced in [13], as an extension of the Hierarchical Dirichlet Process of [14]. It is a useful tool for modelling groups of data coming from different distributions, when we wish to introduce probabilist dependence among them. The HPY assigns to each distribution, G_j , a Pitman-Yor process, [11], all with the same

base measure G_0 . This latter hyperparameter is not fixed by the modeller, but is considered as a random element to be inferred from data. Another Pitman-Yor prior is used as prior for G_0 . This recursive construction has the effect that the support of the G_j s is contained in that of G_0 . As a consequence, all populations share the same random support of G_0 . Denoting with $\{\theta_{j1}, \dots, \theta_{jn_{j..}}\}$ the vector of observations from the j -th population, taking values on a measurable space (Θ, \mathfrak{S}) , the HPY model is described as follows

$$\begin{aligned} \theta_{j,i} | G_j &\sim G_j \quad i = 1, \dots, n_{j..} \\ G_j | d_j, \alpha_j, G_0 &\sim PY(d_j, \alpha_j, G_0) \quad j = 1, \dots, J \\ G_0 | \eta, \gamma, H &\sim PY(\eta, \gamma, H) \end{aligned}$$

where $n_{j..}$ is the number of observations from the j -th population, H is a fixed and diffuse probability measure and the $J+1$ couples of hyperparameters (d_j, α_j) and (η, γ) are chosen to satisfy the conditions $d_j, \eta > 0$, $\alpha_j > -d_j$ and $\eta > -\gamma$, $\forall j \in \{1, \dots, J\}$. Also, the hyperparameters d_j , α_j , η and γ are usually assumed to be unknown and endowed with priors.

The HPY admits the following representation in terms of *Chinese Restaurant Franchise* (See [12] for a detailed description of this culinary metaphor)

$$\theta_{j,i} | \theta_{j,1}, \dots, \theta_{j,i-1}, \alpha_j, d_j, G_0 \sim \sum_{t=1}^{m_{j..}} \frac{n_{jt..} - d_j}{\alpha_j + n_{j..}} \delta_{\theta_{j,t}^*} + \frac{\alpha_j + m_{j..} \cdot d_j}{\alpha_j + n_{j..}} G_0 \quad (1)$$

$$\theta_{j,t}^* | \theta_{1,1}^*, \dots, \theta_{1,m_1}^*, \dots, \theta_{j,t-1}^*, \gamma, \eta, H \sim \sum_{k=1}^K \frac{m_{j,k} - \eta}{\gamma + m_{j..}} \delta_{\theta_k^{**}} + \frac{\gamma + K \cdot \eta}{\gamma + m_{j..}} H \quad (2)$$

where n_{jtk} denotes the number of customers in restaurant j , sitting at table t and trying dish k , while m_{jk} is the number of tables at restaurant j serving dish k , $\{\theta_{j,1}^*, \dots, \theta_{j,m_j}^*\}$ are the dishes served at the m_j tables at restaurant j and K stands for the number of franchise-wide distinct dishes with corresponding labels $\{\theta_1^{**}, \dots, \theta_K^{**}\}$. As in [12], dots in the indexes denote that we are summing over that index, e.g. $m_{j..}$ is the total number of tables in restaurant j .

2.1.1 The missing masses a posteriori and the algorithm

In order to implement the HPY-TS algorithm, the joint posterior of the J missing masses given all observations must be derived. Denoting with $\theta_{\mathbf{n}}$ the joint sample (the array containing observations from all populations), with $\theta_{n_{j..}}$ be the vector of observations from population j and with $A = \{\theta \in \Theta : \theta \notin \theta_{\mathbf{n}}\}$ the set of possible new species, what is needed is the joint distribution of

$$\{G_1(A), \dots, G_J(A)\} | \theta_{\mathbf{n}}, \eta, \gamma, H, d_1, \dots, d_J, \alpha_1, \dots, \alpha_J$$

For easiness of notation, now on we omit the reference to the hyperparameters of the HPY, $\eta, \gamma, H, d_j, \alpha_j$, when conditioning on them. The density of this joint distribution is provided in the following proposition. In the statement, we adopt the notation for tables counts and distinct values previously introduced for the Chinese Franchise Representation of the HPY. Also, $beta(g|a, b)$ stands for a beta density function with parameters a and b , evaluated at g .

Proposition 1. Let $\theta_{\mathbf{n}}$ denote the joined sample from a HPY and let $A = \{\theta \in \Theta : \theta \notin \theta_{\mathbf{n}}\}$. Then, $\{G_1(A), \dots, G_J(A)\} | \theta_{\mathbf{n}}$ admits the following multivariate density

$$f_{\{G_1(A), \dots, G_J(A)\} | \theta_{\mathbf{n}}}(g_1, \dots, g_J) = \int_0^1 \prod_{j=1}^J f_j(g_j | \beta_0, m_{j\cdot}, n_{j\cdot}) \cdot f_0(\beta_0 | K, m_{\cdot\cdot}) d\beta_0$$

where

$$f_j(g_j | \beta_0, m_{j\cdot}, n_{j\cdot}) = beta(g_j | (\alpha_j + m_{j\cdot} d_j) \cdot \beta_0, (\alpha_j + m_{j\cdot} d_j) \cdot (1 - \beta_0) + n_{j\cdot} - d_j \cdot m_{j\cdot})$$

and

$$f_0(\beta_0 | K, m_{\cdot\cdot}) = beta(\beta_0 | \gamma + K\eta, m_{\cdot\cdot} - \eta K)$$

Proof From the Chinese Franchise Represtation of the HPY, the franchise-wide distinct values $\{\theta_1^{**}, \dots, \theta_K^{**}\}$ are governed by G_0 and $G_0 \sim PY(\eta, \gamma, H)$. Applying [10], Corollary 20, the posterior distribution of G_0 , given the observations, satisfies the distributional equation

$$G_0 | \theta_{\mathbf{n}} = \sum_{k=1}^K \beta_k \cdot \delta_{\theta_k^{**}} + \beta_0 \cdot G'_0$$

where

$$\begin{aligned} G'_0 | \theta_{\mathbf{n}} &\sim PY(\eta, \gamma + K\eta, H) \\ \beta | \theta_{\mathbf{n}} = (\beta_0, \dots, \beta_K) | \theta_{\mathbf{n}} &\sim Dir(\gamma + K\eta, m_{\cdot 1} - \eta, \dots, m_{\cdot K} - \eta) \end{aligned}$$

Similarly, we can apply the same result of [10] to G_j to find a distributional equation for G_j , conditionally on G_0 and the data. Also, using the distributional equation for the posterior of G_0 , we find the following distributional equation for G_j

$$G_j | \beta, G'_0, \theta_{\mathbf{n}} m = \sum_{k=1}^K \pi_{j,k} \cdot \delta_{\theta_k^{**}} + \pi_{j,0} \cdot G'_j \quad (3)$$

where

$$\begin{aligned} G'_j | G'_0, \theta_{\mathbf{n}} &\sim PY(d_j, (\alpha_j + m_{j\cdot} d_j) \cdot \beta_0, G'_0) \\ (\pi_{j,0}, \dots, \pi_{j,K}) | \beta, \theta_{\mathbf{n}} &\sim Dir((\alpha_j + m_{j\cdot} d_j) \cdot \beta_0, (\alpha_j + m_{j\cdot} d_j) \cdot \beta_1 + n_{j\cdot 1} - d_j \cdot m_{j1}, \dots \\ &\quad \dots, (\alpha_j + m_{j\cdot} d_j) \cdot \beta_K + n_{j\cdot K} - d_j \cdot m_{jK}) \end{aligned}$$

So, the distribution of $G_j(A) | \theta_{\mathbf{n}}, G_0$ satisfies

$$G_j(A) | \beta, G'_0, \theta_{\mathbf{n}} = \sum_{k=1}^K \pi_{j,k} \cdot \delta_{\theta_k^{**}}(A) + \pi_{j,0} \cdot G'_j(A)$$

$\forall j \in \{1, \dots, J\}$, which implies

$$G_j(A) | \beta_0, \theta_{\mathbf{n}} \sim \text{beta}((\alpha_j + m_j \cdot d_j) \cdot \beta_0, (\alpha_j + m_j \cdot d_j) \cdot (1 - \beta_0) + n_{j\cdot} - d_j \cdot m_j)$$

where we made use of the following facts:

1. $\delta_{\theta_k^{**}}(A) = 0 \forall k = 1, \dots, K$: since $\{\theta_1^{**}, \dots, \theta_K^{**}\} = A^c$.
2. $G'_j(A) = 1$ a.s.: G'_j can be rewritten as $G'_j = \sum_{i \geq 1} \gamma_i \cdot \delta_{\theta_i}$ for some weights $\{\gamma_i\}_{i \geq 1}$ and some atoms $\{\theta_i\}_{i \geq 1}$ i.i.d. from H . Then, $P(\cap_{i \geq 1} \{\theta_i \in A^c\}) = \prod_{i \geq 1} P(\theta_i \in A^c) = \prod_{i \geq 1} 1 = 1$, since H is diffuse. Finally, $P(\cap_{i \geq 1} \{\theta_i \in A^c\}) = 1 \Rightarrow G'_j(A^c) = 0$ a.s., hence $G'_j(A) = 1$ a.s..
3. $\pi_{j,0} | \beta_0, \theta_{\mathbf{n}} \sim \text{beta}((\alpha_j + m_j \cdot d_j) \cdot \beta_0, (\alpha_j + m_j \cdot d_j) \cdot (1 - \beta_0) + n_{j\cdot} - d_j \cdot m_j)$: by the aggregation property of Dirichlet distribution.

Also, since we are conditioning on G_0 (through β, G'_0), $G_j(A) | \beta_0, \theta_{\mathbf{n}}$ is independent of $G_i(A) | \beta_0, \theta_{\mathbf{n}} \forall i, j \in \{1, \dots, J\}, i \neq j$. Hence, their joint distribution is simply the product of the marginals. The last step is to integrate β_0 out

$$\{G_1(A), \dots, G_J(A)\} | \theta_{\mathbf{n}} = \int_0^1 \prod_{j=1}^J G_j(A) | \beta_0, \theta_{\mathbf{n}} \cdot dF_{\beta_0}(\beta_0)$$

where the distribution of β_0 is another beta (again by aggregation of Dirichlet distribution). So, $\{G_1(A), \dots, G_J(A)\} | \theta_{\mathbf{n}}$ admits a density as stated. □

The HPY-TS algorithm can now be easily described. At each time step, it draws a sample from the joint posterior distribution of Proposition 1 and collects the next observation from the the population corresponding to the highest realized value. This strategy outperforms the greedy one that selects the arm with the highest posterior point estimate since it better balances the exploration step. To provide the intuition, suppose to have only a few observations, with an unlucky sample, from a 'winning' arm (a population with a very high species variety), resulting in a low point estimate for its missing mass. This population will not be chosen by the greedy strategy, which only exploits arms with good past behaviour. Whereas, with HPY-TS, the posterior distribution of the missing mass of this population will have high variability, due to the small sample size. This implies a positive probability for that arm to be chosen, if its Thompson draw results in a high value. The HPY-TS is summarized in Algorithm 1.

Algorithm 1: (HPY-TS algorithm)

```

for  $i$  in 1:additional sample do
  draw  $\beta_0 \sim \text{beta}(\gamma + K\eta, m_{..} - \eta K)$ ;
  for  $j$  in 1:J do
    draw  $g_j \sim \text{beta}((\alpha + m_{j,d}) \cdot \beta_0, (\alpha + m_{j,d}) \cdot (1 - \beta_0) + n_{j..} - d \cdot m_{j.})$ ;
  end
  Compute  $j^* = \text{argmax}\{g_j : j \in \{1, \dots, J\}\}$ ;
  Sample the next observation from population  $j^*$ ;
  Update table counts and estimates of the HPY hyperparameters;
end

```

3 Simulated Results

In this section we assess the performances of the HPY-TS algorithm through simulations. We compare it to three alternative strategies

- An *Oracle strategy* that knows the $\{G_1, \dots, G_J\}$ generating the data and always picks the arm with the highest missing mass.
- An *Uniform strategy* that selects the next population uniformly at random.
- A UCB algorithm recently proposed by [2] in the context of security analysis. It is based on the Good and Turing missing mass estimator, derived in [7], and we will refer to it as *Good-Turing strategy*.

We assume as true data generating process J *Zipf laws*, each one supported on a set of size 2500, randomly chosen from a total number of 3000 possible species, hence allowing for a partial sharing of the supports. The parameter of each distribution, $s_j > 1$, controls how the total mass is spread along the support points. If s_j is large, the total mass is concentrated on a few points, while, as s_j approaches 1, the total mass is more spread, with many points of high mass. In the species discovery context, an arm with low parameter s_j is a winning arm.

We consider three different *scenarios*, corresponding to different levels of heterogeneity or homogeneity in species variety across arms. In presence of heterogeneity, a good strategy must be able to detect winning arms soon and play them only. Whereas, in presence of homogeneity, a strategy must be able not to get stuck exploiting only a few arms, but to carefully explore all of them. In our simulations, we fix $J = 8$ and consider the following three scenarios:

1. *Pure Exploitation*, Zipf parameters=(1.3,1.3,2,2,2,2,2,2): in this scenario, there are two 'winning' arms. A good strategy should be able to intensively exploit these two arms, without exploring much the other six suboptimal arms.
2. *Pure Exploration*, Zipf parameters=(1.3,1.3,1.3,1.3,1.3,1.3,2,2): in this scenario, the majority of arms are equally profitable. A good strategy should not get stuck exploiting just a few of them, but continue to explore all the six good arms.

3. *Exploration plus Exploitation*, Zipf parameters=(1.3,1.3,1.3,1.3,2,2,2,2): in this scenario, there are four good arms and four bad ones. A good strategy should adequately balance exploitation and exploration, by stopping to play the four suboptimal arms soon, but continuing to play all the other four.

Figure 1 reports the results of simulations in these three scenarios. Each panel in Fig. 1 displays the average number of species discovered by the four algorithms, as a function of the additional observed samples. In particular, the results are averages of 60 runs. For each run, we assume an initial sample of 30 observations per arm to be available and collect further 300 observations, following the four possible strategies. The hyperparameters of the HPY are endowed with priors, $\eta, d_1, \dots, d_J \sim \beta(1, 2)$ i.i.d. and $\gamma, \alpha_1, \dots, \alpha_J \sim \exp(1)$ i.i.d..

In the simulations, HPY-TS algorithm performs well in all scenarios, discovering less new species than the Oracle strategy, but more than the Uniform and to the Good-Turing strategy. Figure 1 shows how these latter strategies seem to balance the exploration-exploitation trade-off worse than HPY-TS. They perform relatively well only in the two extreme cases of pure exploration or pure exploitation, charts (a) and (b) in Fig. 1. On the one hand, the Good-Turing strategy does too much exploitation. It focuses on exploiting just a few arms (those with good past performances), without continuing to explore the others. In fact, it performs quite well in the pure exploitation scenario, Fig. 1 (a), but quite poorly, in presence of more 'winning' arms, panel (b) and (c) of Fig. 1. On the other hand, as expected, the Uniform strategy does too much exploration. It continues to play all arms, irrespectively of their past behaviours. Its performances are very poor, except in the extreme scenario of pure exploration, Fig. 1 (b). Instead, the HPY-TS algorithm seems to be robust to changes in species variety across arms, without the need of any tuning parameter to regulate the exploration and exploitation levels. In all scenarios, it performs well, standing behind only to the Oracle strategy. In particular, in the intermediate scenario, Fig. 1 (c), its results are very close to the Oracle's ones, while in the extreme cases, Fig. 1 (a) and (b), it is still as good as or better than both the Uniform and the Good-Turing strategies.




Fig. 1 (a) Pure Exploitation scenario; (b) Pure Exploration scenario; (c) Exploration and Exploitation Scenario.

References

- [1] Auer, P., Cesa-Bianchi, N. and Fischer, P.: Finite-time analysis of the multiarmed bandit problem. *Machine Learning* **47**(2), 235–256 (2002)
- [2] Bubeck, S., Ernst, D. and Garivier, A.: Optimal discovery with probabilistic expert advice: finite time analysis and macroscopic optimality. *JMLR* **14**, 601–623 (2013)
- [3] Bunge, J. and Fitzpatrick, M.: Estimating the number of species: a review. *JASA* **88**, 364–373 (1993)
- [4] Favaro, S., Lijoi, A., Mena, R.H. and Pruenster, I.: Bayesian nonparametric inference for species variety with a two parameter PoissonDirichlet process prior. *JRSSB* **71**, 993–1008 (2009)
- [5] Favaro, S., Lijoi, A. and Pruenster, I.: A new estimator of the discovery probability. *Biometrics* **68**, 1188–1196 (2012)
- [6] Favaro, S., Lijoi, A. and Pruenster, I.: Asymptotics for a Bayesian nonparametric estimator of species variety. *Bernoulli* **18**(4), 1267–1283 (2012)
- [7] Good, I.J.: The population frequencies of species and the estimation of population parameters. *Biometrika* **40**, 237–64 (1953)
- [8] Lai, T.L. and Robbins, H.: Asymptotically efficient adaptive allocation rules. *Adv. App. Math.* **6**, 4–22 (1985)
- [9] Lijoi, A., Mena, R.H. and Pruenster, I.: Bayesian nonparametric estimation of the probability of discovering a new species. *Biometrika* **94**, 769–786 (2007)
- [10] Pitman, J.: Some developments of the Blackwell-MacQueen urn scheme. In: *Statistics, Probability and Game Theory* (eds. T. S. Ferguson, L. S. Shapley and J. B. MacQueen), pp. 245–267. Institute of Mathematical Statistics, Hayward. (1996)
- [11] Pitman, J. and Yor, M.: The two parameter Poisson-Dirichlet distribution derived from a stable subordinator. *Ann. of Prob.* **25**, 855–900 (1997)
- [12] Teh, Y.W. and Jordan, M.: Hierarchical Bayesian Nonparametric Models with Applications. In: *Bayesian Nonparametrics* (eds. N.L. Hjort, C. Holmes, P. Mueller and S.G. Walker), pp. 158–207. Cambridge University Press. (2010)
- [13] Teh, Y. W.: A hierarchical Bayesian language model based on PitmanYor processes. In: *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, 98592. Association for Computational Linguistics, Morristown. (2006)
- [14] Teh, Y.W., Jordan, M., Beal, M. and Blei, D.: Hierarchical Dirichlet processes. *JASA* **101**, 1566–81 (2006)
- [15] Thompson, W.R.: On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* **25**, 285–294 (1933)