Published in final edited form as:

Pharm Stat. 2017 November ; 16(6): 414–423. doi:10.1002/pst.1819.

Bayesian Nonparametric Statistics: A New Toolkit for Discovery in Cancer Research

Peter F. Thall¹, Peter Mueller², Yanxun Xu³, and Michele Guindani⁴

¹Department of Biostatistics, University of Texas, M. D. Anderson Cancer Center

²Department of Mathematics, University of Texas at Austin

³Department of Applied Mathematics and Statistics, Johns Hopkins University

⁴Department of Statistics, University of California, Irvine

Abstract

Many commonly used statistical methods for data analysis or clinical trial design rely on incorrect assumptions, or assume an over-simplified framework that ignores important information. Such statistical practices may lead to incorrect conclusions about treatment effects, or clinical trial designs that are impractical or that do not accurately reflect the investigator's goals. Bayesian nonparametric (BNP) models and methods are a very flexible new class of statistical tools that can overcome such limitations. This is because BNP models can accurately approximate any distribution or function, and can accommodate a broad range of statistical problems, including density estimation, regression, survival analysis, graphical modeling, neural networks, classification, clustering, population models, forecasting and prediction, spatio-temporal models, and causal inference. This paper describes three illustrative applications of BNP methods, including a randomized clinical trial to compare treatments for intraoperative air leaks after pulmonary resection, estimating survival time with different multi-stage chemotherapy regimes for acute leukemia, and evaluating joint effects of targeted treatment and an intermediate biological outcome on progression-free survival time in prostate cancer.

Keywords

Bayesian nonparametric statistics; Clinical trial design; Density estimation; Dynamic treatment regime; Targeted Therapy

1. INTRODUCTION

Many statistical methods commonly used for analyzing medical data or designing clinical trials rely on incorrect assumptions, or they assume an over-simplified framework that ignores important information. A common example is a statistical analysis based on the assumption that the data are normally distributed, despite the fact that the histogram is not bell-shaped or has multiple modes. A second example is comparison of survival times for

frontline anti-cancer treatments while ignoring subsequent salvage therapies, which fails to account for effects of salvage on survival. A third example is using the Cox model[1] to analyze survival data from two treatments having Kaplan-Meier plots[2] that cross, which implies that the proportional hazards assumption underlying the Cox model cannot be correct. Such statistical practices may lead to incorrect conclusions about treatment effects, which serve physicians and patients poorly.

In these and other examples, analysis of medical data often involves an unknown probability distribution, which we denote by G . Statistical inference usually proceeds by assuming that G is a member of some parametric family, such as normal or Weibull families of distributions, having unknown parameters. Similar assumptions are made for an unknown function, f , which may describe how the probability of response varies with dose, or may be the longitudinal profile of a biological variable, such as prostate specific antigen, that is recorded repeatedly over time. Key questions are how the distribution G or function f depend on treatment or other patient prognostic covariates. If one assumes a particular parametric form for G or f , and this assumed model is incorrect or overly restrictive, inference will be restricted to that assumed form, possibly leading to incorrect conclusions. In some applications, the conventional statistical practice of fitting several models to a dataset and choosing the model that gives the best fit may mitigate this problem, but it still is limited by the flexibility of models that are considered.

Bayesian nonparametric (BNP) models and methods (see, for example, Hjort, et al.[3], Walker[4] or Mueller, et al.[5]) are a new class of statistical tools that can overcome some limitations of conventional methods, including the specific issues in the examples described above. BNP models are very flexible, and the family of BNP models can accommodate a broad range of different data structures and statistical problems. These include density estimation, regression, survival analysis, graphical modeling, neural networks, classification, clustering, population models, forecasting and prediction, spatio-temporal models, and causal inference (Mitra and Mueller[6]). A major advantage of BNP models is that, based on the observed data, they can accurately approximate essentially any distribution or function, a property known as “full support.” For example, Barrientos, et al.[7] give a formal explanation of such full support properties for the widely used dependent Dirichlet process model (MacEachern[8]; De Iorio, et al.[9]). For this reason, the class of BNP probability models and accompanying statistical methods can, in some cases, help investigators avoid making invalid statistical inferences due to assuming an overly restrictive model. An attractive feature of BNP models is that they often identify unexpected structures in a dataset that cannot be seen using conventional statistical models and methods. Examples of such structures include patient clusters, treatment-subgroup interactions, gene signatures, and complex patterns of biomarker change over time.

This paper presents three examples that illustrate how BNP models and methods can be applied to obtain insights that may not be provided by conventional statistical methods. The examples include a clinical trial to compare treatments for intraoperative air leaks (IALs) after pulmonary resection, estimation of expected survival time with different multi-stage chemotherapy (chemo) regimes for acute leukemia, and evaluation of the joint effects of a targeted treatment and an intermediate biological outcome on progression-free survival

(PFS) time in prostate cancer. We focus on describing how each BNP method addresses critical aspects of the problem that conventional statistical models and methods may not accommodate. In describing each analysis, we minimize technical details, include graphical illustrations, and provide references to methodological papers. Recent reviews of BNP models and methods are given, for example, by Mueller et al.[5], and Dunson[10], and specifically for biomedical problems by Mitra and Mueller [6].

2. COMPARING TREATMENTS FOR AIR LEAKS AFTER LUNG SURGERY

Xu, et al. [11] apply a BNP model to construct a clinical trial design for patients who experience intraoperative air leaks (IALs) after pulmonary resection. The goal of the trial is to compare the use of a novel hydrogel sealant, Progel, to the standard procedure of using sutures and staples. Although Progel has been approved by the U.S. FDA, its efficacy for reducing IALs in lung resection patients has not been established, and it is not used routinely.

Historical data from 46 mesothelioma patients who underwent lung resection and received standard care for IALs show that the distribution of Y = time to resolve an air leak is not bell-shaped and does not follow any standard parametric form, like a normal, lognormal, or Weibull distribution. Figure 1 (left) shows that the histogram of Y values has an irregular shape, with multiple modes. The trial's hypothesis is that Progel may reduce the average time to resolve an IAL, and it possibly may prevent IALs in some patients ($Y=0$). The multimodality and possibility that $Y=0$ present technical complications that standard statistical distributions and methods cannot accommodate. To construct a clinical trial design to compare Progel to standard care, conventional methods that compare mean times of the two treatments based on a parametric distribution for computing power and sample size are not appropriate. First, a conventional test generally ignores the spike at 0. Given the irregular shape of the distribution in Figure 1 (left), and the possibility that $Y=0$ with Progel, summarizing the two IAL distributions in terms of their means is misleading.

Conventional approaches also lead to designs requiring sample sizes that were not considered feasible in the single-institution setting where the trial was planned. For example, assuming that Y is normal, based on the historical mean of 8 days and standard deviation 8.76, a two-sample one-sided .05-level t-test with power = 0.80 to detect a 25% drop in the mean, from 8 to 6 days, would require a sample of 476 patients.

Alternatively, using a log scale one could compute sample size based on a historical mean of 1.61 and standard deviation 0.97, and a drop from $\log(8)$ to $\log(6)$ days. This design would require a sample size of 280 patients, which still would be an infeasible accrual target for this medical condition and treatment in a single institution trial.

For later comparison, we briefly outline how a conventional parametric analysis that accounts for the possibility that $Y=0$ might be set up. As a conventional comparator, we can assume a zero-enriched Weibull distribution, that is, a two-component mixture of a point mass at 0 and a Weibull. Inference under this model can be formulated easily based on maximum likelihood estimation. The main limitation of this setup is that actual differences

between the distributions under Progel and standard care might not be well characterized by shifts of the entire distribution, which is all that the Weibull model can accommodate.

Taking a BNP approach to constructing a design, the first step is to apply BNP density estimation (Mueller and Mitra [12]), to estimate the distribution, denoted by G_0 , of the IAL resolution times for standard care based on the historical data. This estimate is used as a basis for comparison, and for constructing a hypothesized distribution, G_1 , for Progel. The BNP estimate of G_0 and hypothesized G_1 are given in Figure 1 (middle), which shows that G_0 is a smoothed version of the historical data histogram (black curve), and G_1 (red dashed curve) is a left-shifted version of G_0 . The red spike at 0 for G_1 formalizes the hypothesis that air leaks might not develop at all in some patients given Progel, although the model is formulated generally to allow the possibility that the standard distribution, G_0 , also may have a spike at 0. A major advantage of BNP density estimation is that it naturally accommodates irregular shaped distributions, including the spike at 0.

Inference, based on the trial data, comparing G_1 to G_0 under the BNP model characterizes uncertainty about the fitted distribution G_0 , rather than treating it as fixed. The estimated density of G_0 in Figure 1 (middle) illustrates how the BNP model, known as a dependent Dirichlet process (DDP) (MacEachern, 1999; De Iorio, et al, 2004), can be applied to estimate the distribution of IAL resolution times based on the historical data. The DDP estimate of G_0 is essentially a weighted average, or “mixture,” of normal distributions, with mixture components corresponding to the peaks in the histogram in Figure 1 (left). The joint model for the two distributions requires G_1 to be shifted to the left of G_0 , formally, they are “stochastically ordered,” and it also allows G_1 or G_0 to have a spike at 0, corresponding to immediate resolution of air leaks. Although this did not occur with standard care in the historical data, it is clinically possible and is accommodated by the probability model. Therefore, we assume that each of the distributions G_1 (Progel) and G_0 (standard care) of time to resolve an IAL is an average (mixture) of a point mass at 0 (no air leak) and a weighted sum of log normal distributions having different means. These BNP distributions are extremely flexible, and can closely fit virtually any observed data, which in turn leads to a very reliable clinical trial design.

The stochastic ordering in the joint model of G_1 and G_0 formalizes the assumptions that, compared to the standard, Progel is more likely to have no air leak, and also has shorter expected resolution times for any air leaks that occur. It also implies that G_1 has smaller means than G_0 for the normal distributions in the weighted sum. These model assumptions reflect the facts that Progel is inert, and thus cannot react chemically with the patient's lung tissue, is not a potential source of infection, does not slow down the healing process, and does not contribute to air leak formation.

The trial design based on the BNP model uses decision rules based on utilities of possible Y values, described below, that were elicited from the trial's Principal Investigator (PI). This design requires a maximum of 48 patients, much smaller than 476 or 280 patients required by conventional methods. In simulations, we found this to be a sufficient sample size for meaningful inference under the proposed BNP model under a set of realistic scenarios.

To quantify the clinical relevance of each possible IAL resolution time Y , numerical utilities quantifying the desirabilities of the possible values of Y were elicited from the oncologist planning the trial. These are shown in Table 1, and plotted in the right side of Figure 1. The numerical utility, $u(Y)$, of each resolution time Y quantifies its relative desirability.

The highly non-linear nature of the clinical preferences in $u(Y)$ reflects the large clinical utility of early resolution of IALs, with short resolution times far more preferable than larger times beyond 10 days. An average reduction of 1 day is far more important for short resolution times, e.g. from 6 to 5 days, than the same reduction from a much larger resolution time, say from 30 to 29 days. The values of $u(Y)$ show why any conventional two-sample test based on comparing the means of Y is inappropriate. Aside from these issues, the sample sizes of 476 or 280 required by conventional t-tests to compare means are far too large for running a single-institution a trial in a practicable time frame.

Given the general BNP model setup, the inferential targets used by the comparative test are the mean utilities, denoted by U_1 and U_0 . These are computed with respect to G_1 and G_0 based on the observed data. Below, we describe a group sequential test that concludes Progel is superior to standard care if $\Pr(U_1 > U_0 + 18 \mid \text{data})$ is sufficiently large. This says that, given the observed data, it is likely that Progel provides a mean advantage over standard care of at least 18 points on the utility scale. The targeted improvement 18 for the mean utility was elicited from the PI by showing him several different desirable distributions G_1 for Progel and their corresponding mean utilities.

During the trial, this posterior probability is compared to test cut-offs after cohorts of 16, 32, and a maximum of 48 patients, with the test cut-offs derived to control the overall false positive probability at .05. In the protocol design, a threshold of 90% was used for stopping for efficacy and 5% for stopping for futility.

Extensive simulations, given in Tables 2 and 3 of Xu, et al. [11], showed that this design has high power to detect true mean utility differences under a wide range of possible cases where Progel provides an actual benefit over the control. Since the advantage of the design over a design based on a two-sample t-test may be ascribed to both the BNP model and the use of $u(Y)$ in place of Y , a design using $u(Y)$ but based on a zero-enriched Weibull model also was studied. Simulations for comparing the designs based on the BNP model and the zero-enriched Weibull model, which were presented in Table 4 of Xu, et al. [11], showed that the BNP model-based design compares quite favorably, and has much higher power and smaller type I error. In five simulated scenarios, the difference in probabilities of correct decision for the BNP design and conventional designs ranged from .09 to .46. This large advantage is due to the fact that the BNP model can accommodate, for example, increased probabilities of short resolution time. The expected utility criterion is very sensitive to such changes, since early days are given much larger utility weights than later days.

3. ADJUSTING FOR LACK OF RANDOMIZATION

Estey, et al. [13] reported results of a four-arm randomized clinical trial designed to compare the effects of remission induction chemotherapy (chemo) combinations on 210 patients with

acute myelogenous leukemia (AML) or myelodysplastic syndrome (MDS). Patients with AML or MDS have blast cell percentages far above the normal range of 3 – 5%, and low platelet and white cell counts. Thus, the first goal of induction chemo for AML/MDS is to achieve a complete remission (CR), which means that these three types of blood cells have returned to normal functional ranges. CR is defined as the patient having less than 5% blast cells, a platelet count of at least $10^5/\text{mm}^3$, and a white blood cell count of at least $10^3/\text{mm}^3$. CR is a necessary but not sufficient condition for long term survival, however, since patients who achieve CR may relapse or die.

Denoting FAI = fludarabine + ara-C + idarubicin, GCSF = granulocyte colony stimulating factor, ATRA = all-trans retinoic acid, and HDAC = high dose ara-C, the four chemos were FAI, FAI+ATRA, FAI+GCSF, and FAI+ATRA+ GCSF. The goal was to assess the effects on patient outcome of adding ATRA, G-CSF, or both to the baseline chemo FAI. Based on conventional Cox model, logistic regression, and logrank test analyses, Estey et al. [13] concluded that G-CSF had a beneficial effect on the probability of CR, but that neither G-CSF nor ATRA had any significant effects on either overall survival or event-free survival among patients who achieved CR.

To account for additional effects of salvage treatments on overall survival (OS) time, Wahed and Thall [14] re-analyzed this trial's data. Salvage chemo was given at two different possible stages, if the patient's disease was resistant to induction (salvage1), or at disease progression following complete remission (CR) achieved by induction (salvage2). The possible combinations of treatment sequences and patient outcomes are illustrated in Figure 2. The numbers of patients with outcomes of each type following induction, and following salvage therapy, are given in Table 2.

A dynamic treatment regime (DTR) in the trial consisted of the triple (frontline, salvage1, salvage2). Categorizing salvage by whether or not it contained high dose ara-C (HDAC), there were a total of 16 possible DTRs. While all patients received one of the four induction chemo combinations, depending on their induction outcome (CR, disease resistant to induction, died during induction) each patient could receive either salvage1 or salvage2, but not both. The flow diagram in Figure 2 shows how each patient's overall survival time was one of four possible sums, (1) time to death during induction chemo, (2) time to achieve CR + time to death in CR, (3) time to CR + time to progression + subsequent survival time (4) time to resistance + subsequent survival time.

When accounting for the DTRs, a key problem is that patients were not randomized between different types of salvage chemos, which may introduce selection bias into the estimates. Wahed and Thall[14] dealt with this by using inverse probability of treatment weighting (IPTW, (Murphy et al., [16], van der Laan and Petersen, [17], Robins et al.[18]). Wahed and Thall [14] focused on estimating mean OS for each DTR, using both IPTW and likelihood based methods. Their results, summarized in Table 6 of Wahed and Thall[14], showed that the two methods gave numerically very different estimates, but they gave the same conclusions regarding the comparative effects of the 16 DTRs. They concluded that mean OS time was smallest for the four strategies with FAI alone as frontline regardless of salvage therapies, with the two exceptions that (FAI+GCSF+ATRA, HDAC, HDAC) was

slightly inferior to (FAI, Other, HDAC) and (FAI, Other, Other). They also concluded that FAI+ATRA was the best remission induction therapy, the type of salvage therapy was irrelevant for patients with disease resistant to FAI+ATRA, and HDAC was a superior salvage therapy for patients who achieved CR with FAI+ATRA but later relapsed.

Xu, et al.[15] re-analyzed this dataset, with the aim to obtain improved estimates of mean OS time by exploiting the flexibility of BNP models. Inference using the more general augmented IPTW (AIPTW; Tsiatis [19], Moodie, et al. [20], Zhao, et al.[21]) hinges on the correctness of the assumed model or, in the case of AIPTW, the correctness of at least one of the assumed models for either the propensity score (treatment assignment probability) or observed outcome. Reliable inference is not possible when neither model is known. BNP inference can fill this gap, because BNP models always provide a good fit to the data, due to the fact that the prior model has full support. Xu, et al. [15] applied a DDP with a Gaussian process prior (DDP-GP model), to account for possible patient covariate effects on treatment selection at any stage of the DTR.

While we briefly summarize the BNP model, these technical details are not needed in the following discussions and can safely be skipped, unless a reader is interested in the specific model. The diagram in Figure 2 includes 7 possible transition times, including, for example, T^c for the transition from the initial state to CR. We constructed a BNP survival regression model by modeling each possible transition time based on the patient's current history, including baseline covariates and preceding transition times. Let T^k denote the k -th transition time in the flow chart of Figure 2, for $k = 1, \dots, 7$. The BNP model assumes that each $\log(T^k)$ follows a DDP mixture (similar to the model in §2), with the mean of each normal distribution in the mixture following a Gaussian process (GP) prior that is a function of the patient's history up to the k th transition, the DDP-GP model noted earlier.

The overall BNP model is constructed from seven such DDP-GP models, one for each possible transition time. The relevant summary is OS time, which is determined by sums of transition probabilities along four possible paths through the diagram. A patient's OS time (summarized in Figure 2) could be (i) the time to death during induction, (ii) the time to achieve CR plus the time to death in CR, (iii) the time to resistance plus the subsequent time to death following salvage, or (iv) the sum of the times to CR, progression after CR, and death after progression. The BNP model flexibly accounts for these four possibilities, covariate effects, and the relevant frontline and salvage treatment effects on each transition time. See Xu, et al. [15] for technical details.

To validate the DDP-GP model-based approach, Xu et al. [15] performed an extensive simulation study in which treatments were chosen based on the values of patient covariates, as a physician would do in practice. Thus, selection bias was built into the simulation and its magnitude was known. Figure 3 (left) shows the results for eight two-stage regimes, based on the simulated data, comparing the BNP method, IPTW, and AIPTW. In Figure 3 (left), the vertical axis gives the difference between each estimated posterior mean OS time and the "true" mean used in the simulations, so values closer to 0 are more desirable. Each case was simulated 1,000 times. In each notched box-whisker plot, the box gives the interquartile range (IQR) from the 25th percentile (Q1) to 75th percentile (Q3), the mid-line is the median,

the top whisker limit is $Q3 + 1.5 \text{ IQR}$, and the bottom whisker limit is $Q1 - 1.5 \text{ IQR}$. The plots show that, across eight different strategies, the BNP method (yellow boxes) gives estimates that are both more accurate and more reliable than those provided by the two conventional methods (green and blue boxes) for bias correction. This may be attributed to the fact that the complex simulated survival time distributions were fit very accurately by the BNP model.

Figure 3 (right) presents similar box plots, based on the data reported by Estey, et al. [13], for the regimes indexed by 1 = (FAI, HDAC, HDAC), 2 = (FAI+ATRA, HDAC, OTHER), 3 = (FAI+GCSF, HDAC, HDAC), 4 = (FAI+ATRA+GCSF, OTHER, OTHER). Covariate-specific posterior mean OS estimates obtained by the BNP-model-based method are given for four combinations of patient age and cytogenetic abnormality. While the posterior estimates have substantial variability, they suggest that, for younger patients, regime 4 is best for poor cytogenetics and regime 3 is best for intermediate cytogenetics. They also suggest that, for older patients, regime 1 is best for either cytogenetic subgroup. While, unfortunately, no regime provides a substantive improvement in OS, this illustrates how BNP model-based methods can be used to optimize personalized treatments or multi-stage regimes.

4. ESTIMATING TARGETED AGENT EFFECTS ON SURVIVAL

Mathew et al. [22] reported results of a randomized trial of Docetaxel + Imatinib (D+I) versus Docetaxel (D) in men with advanced prostate cancer. The trial's primary goal was to assess the effect on progression-free survival (PFS) time of adding Imatinib to Docetaxel. It was hypothesized that Imatinib would reduce the concentration of phosphorylated platelet-derived growth factor (p-PDGFR) in the blood, and that this in turn would inhibit tumor angiogenesis and reduce the incidence of bone metastases, and thus improve PFS. A total of 88 patients were enrolled in the study (47 in the D arm and 41 in the D+I arm). For each patient, p-PDGFR was measured for each of a large number of cells in each of two blood samples, one taken pre-chemo and the other post-chemo. A key parameter was the effect of the within-patient change, $C = C_{\text{post}} - C_{\text{pre}} = [\text{post-chemo p-PDGFR concentration}] - [\text{pre-chemo p-PDGFR concentration}]$, on PFS. The motivating ideas were that, on average, smaller (or ideally, negative) values of C , corresponding to a drop in p-PDGFR after chemo, would be associated with longer PFS, and that D+I would yield smaller values of C compared to D alone. Assessing this required accurate estimates of the distributions of C_{post} and C_{pre} for each patient.

It might seem that simply computing the means of C_{post} and C_{pre} would be sufficient. However, histograms of these variables based on the within-patient blood cell samples show that the distributions have multiple modes in some patients. Figure 4 gives histograms of C_{post} (pink) and C_{pre} (grey) cell measurements for three patients. The bimodality implies that reducing each sample to its mean misrepresents the C_{pre} and C_{post} values.

To assess association of change from C_{pre} to C_{post} with treatment and PFS, a useful parameter is $\Pr(C_{\text{pre}} < C_{\text{post}})$, the probability that the patient's p-PDGFR went up after chemo. A value $> .50$ corresponds to p-PDGFR going up, $< .50$ to p-PDGFR going

down, and $\beta = .50$ to no change in p-PDGFR. Mathew et al. [22] estimated β very reliably for each patient from the large numbers of cells in each blood sample. A conventional parametric set up may consider a Cox model for PFS, but this was shown to provide a very poor fit to this data set. Alternative parametric time-to-event regression models still rely on a single point estimate of β , for use as a covariate, which does not take into account the complexity and multimodality of the p-PDGFR samples (see Mathew et al. [22]).

Motivated by this problem, Graziani, et al. [23] reported a BNP analysis of the trial data. Figure 4 gives BNP density estimates, plotted as pink dashed lines for C_{post} and solid black lines for C_{pre} , corresponding to each histogram. These density estimates were obtained by assuming BNP models (Dirichlet process mixtures of normals) for the distributions of C_{post} and C_{pre} . As shown by Figure 4, the BNP density estimates identified the multiple modes in the distributions of C_{post} and C_{pre} for each patient, allowing the numbers and locations of the modes to vary from patient to patient.

The BNP formulation used a survival time regression model to quantify the effects of treatment on the levels of the p-PDGFR values in terms of β . The model also provided a framework for estimating how interactions between β and treatment may affect PFS, formalized in equation (7) of Graziani, et al. [23]. The BNP analysis identified clusters of values, shown in Figure 5. This figure illustrates that, contrary to the hypothesis of Mathew, et al. [22], p-PDGFR went down more often for patients in the D+I arm. Also, contrary to the hypothesis, the BNP analysis of the joint effects of treatment, β , and prognostic covariates on PFS showed that an increase in p-PDGFR was associated with longer rather than shorter PFS. This was the case both overall (Figure 6, left side) and predictively for individual patients (Figure 6, right side). After accounting for these effects, there was little difference in PFS between the D+I and D arms.

5. CONCLUSIONS

The illustrations show how BNP models and methods may be applied in settings where conventional statistical methods are inadequate or inappropriate. For the intra-operative air leak clinical trial, BNP density estimation accounted for the complex multi-modal structure of the historical resolution time distribution. This provided a basis for stating hypotheses in terms of entire distributions, rather than in terms of means based on unimodal distributions, which clearly were inappropriate. By exploiting elicited utilities and this density estimate, the BNP approach accounted for the fact that a given reduction in resolution time had different desirabilities for large versus small resolution times. This gave a design with a practical sample size, while conventional approaches did not.

For analysis of DTRs in chemotherapy of acute leukemia, the BNP model provided superior estimation of overall mean survival time, and covariate-specific survival estimates to identify optimal individualized multi-stage treatment regimes. For the prostate cancer trial data, the BNP model facilitated cluster identification and inferences about key relationships between targeted treatment, biomarker change, and PFS time. Methodological and technical details and online tools for implementation of the three examples that we have reviewed are given in the original publications.

While BNP models often are mathematically complex, their properties are intuitively clear in the sense that distributions are represented as mixtures. Consequently, BNP inference is usually more robust against moderate changes in the data and outliers, because the mixture structure of typical BNP models explicitly account for heterogeneity and irregular patterns in the distribution. Quantifying such uncertainties is important even with small or moderately large datasets, as illustrated by the clinical trial design example.

Finally, some words of caution are needed. BNP inference is no panacea. Implementation requires special purpose computer software. Such software is freely available (see, for example, Jara [24]), but it requires statistical expertise for proper model construction. BNP methods are most advisable for settings where the scientific inference problem calls for the level of generality and flexibility that they provide.

Acknowledgments

This work was supported by NCI Cancer Center Core Grant CA016672. Peter Thall's research was supported by NCI grant RO1 CA 83932. Peter Mueller and Yanxun Xu's research was supported by NCI grant RO1 R01 CA157458-01A1. The authors thank two reviewers for their constructive comments.

References

1. Cox DR. Regression models and life-tables, *Journal of the Royal Statistical Society. Series B (Methodological)*. 1972; 34(2):187–220.
2. Kaplan EL, Meier P. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*. 1958; 53(282):457–481.
3. Hjort, NL.Holmes, C.Mueller, P., Walker, SG., editors. *Bayesian Nonparametrics*. Cambridge University Press; 2010.
4. Walker, S. Bayesian nonparametrics. In: Damien, P.Dellaportas, P.Polson, NG., Stephens, DA., editors. *Bayesian Theory and Applications*. Oxford University Press; 2013. p. 249-270.
5. Mueller, P., Quintana, F., Jara, A., Hanson, T. *Bayesian Nonparametric Data Analysis*. Springer-Verlag; 2015.
6. Mitra, R., Mueller, P., editors. *Nonparametric Bayesian Inference in Biostatistics*. Springer-Verlag; 2015.
7. Barrientos AF, Jara A, Quintana FA. On the support of MacEachern's dependent Dirichlet processes and extensions. *Bayesian Anal*. 2012; 7(2):277–310.
8. MacEachern, S. Dependent nonparametric processes, *ASA Proceedings of the Section on Bayesian Statistical Science*. American Statistical Association; Alexandria, VA: 1999.
9. De Iorio M, Mueller P, Rosner GL, MacEachern SN. An ANOVA Model for Dependent Random Measures. *Journal of the American Statistical Association*. 2004:205–125.
10. Dunson D. Nonparametric Bayes applications to biostatistics, in Hjort et al. 2010:223–273. 2010.
11. Xu Y, Thall PF, Mueller P, Mehran RJ. A Bayesian nonparametric utility-based design for comparing treatments to resolve air leaks after lung surgery. *ArXiv e-prints*. 2016a; 1506:07687.
12. Mueller P, Mitra R. Bayesian nonparametric inference why and how. *Bayesian Anal*. 2013; 8(2): 269–302.
13. Estey EH, Thall PF, Pierce S, Cortes J, Beran M, Kantarjian H, Keating MJ, Andreeff M, Freireich E. Randomized phase II study of Fludarabine +Cytosine Arabinoside+ Idarubicin +/- All Trans Retinoic Acid +/- Granulocyte-colony stimulating factor in poor prognosis newly diagnosed acute myeloid leukemia and myelodysplastic syndrome. *Blood*. 1999; 93:2478–2484. [PubMed: 10194425]
14. Wahed AS, Thall PF. Evaluating joint effects of induction-salvage treatment regimes on overall survival in acute leukemia. *J Royal Statistical Society, Series C (Applied Statistics)*. 2013; 62:67–83.

15. Xu Y, Mueller P, Wahed A, Thall P. Bayesian nonparametric estimation for dynamic treatment regimes with sequential transition times. *Journal of the American Statistical Association*. 2016b; 921–950. [PubMed: 28018015]
16. Murphy S, Van Der Laan M, Robins J. Marginal mean models for dynamic regimes. *Journal of the American Statistical Association*. 2001; 96(456):1410–1423. [PubMed: 20019887]
17. van der Laan MJ, Petersen ML. Causal effect models for realistic individualized treatment and intention to treat rules. *International Journal of Biostatistics*. 2007; 3(1)
18. Robins J, Orellana L, Rotnitzky A. Estimation and extrapolation of optimal treatment and testing strategies. *Statistics in medicine*. 2008; 27(23):4678–4721. [PubMed: 18646286]
19. Tsiatis, A. *Semiparametric theory and missing data*. Springer: 2007.
20. Moodie EE, Richardson TS, Stephens DA. Demystifying optimal dynamic treatment regimes. *Biometrics*. 2007; 63(2):447–455. [PubMed: 17688497]
21. Zhao Y-Q, Zeng D, Laber EB, Song R, Yuan M, Kosorok MR. Doubly robust learning for estimating individualized treatment with censored data. *Biometrika*. 2015; 102:151–168. [PubMed: 25937641]
22. Mathew P, Thall PF, Jones D, Perez C, Bucana C, Troncoso P, Kim S-J, Millikan R, Fidler IJ, Logothetis C. Platelet-derived growth factor receptor inhibitor imatinib mesylate and docetaxel: A modular phase I trial in androgen-independent prostate cancer. *J Clinical Oncology*. 2004; 22:3323–3329.
23. Graziani R, Guindani M, Thall PF. Bayesian nonparametric estimation of targeted agent effects on biomarker change to predict clinical outcome. *Biometrics*. 2015; 71:188–197. [PubMed: 25319212]
24. Jara A. Applied Bayesian non-and semi-parametric inference using DPpackage. *Rnews*. 2007; 7:17–26.

Figure 1.

Left side: Time to resolution of intraoperative air leaks (historical data). Middle: Density estimates for time to resolution of intraoperative air leak (IAL) under control (G_0 , black), and hypothesized density under Progel (G_1 , red). Right Side: Plot of elicited utility as a function of time to resolve an IAL.

$$OS = \begin{cases} T^D, & \text{if death during induction} \\ T^R + T^{RD}, & \text{if death after salvage for resistant disease} \\ T^C + T^{CP} + T^{PD}, & \text{if death after salvage for progression after CR} \\ T^C + T^{CD}, & \text{if death in CR} \end{cases}$$

Figure 2.

Flow diagram showing the possible combinations of treatment (ellipses, with lower case tags) and patient outcomes (rectangles, with upper case tags) in the trial comparing chemotherapies for AML.

Figure 3.

Left side: Box plots of estimated mean OS for the BNP, IPTW, and A-IPTW methods, based on simulated data for eight strategies with treatments chosen based on patient covariates.

Right side: Box plots of estimated covariate-specific mean OS, based on the fitted BNP model for the AML data, for four strategies and four combinations of age and cytogenetics.

Figure 4. Histograms of the p-PDGFR values C_{post} (pink) and C_{pre} (gray) for three patients, along with the corresponding BNP density estimates of the distributions.

Figure 5. Clusters of posterior means of $\Pr(C_{\text{pre}} < C_{\text{post}})$, the probability of an increase in p-PDGFR values from pre-chemotherapy to post-chemotherapy.

Figure 6.

Left side: Kaplan Meier estimates of PFS for patients experiencing a drop (dashed line) or increase (solid line) in p-PDFGR. Right side: Median of the distribution of the posterior predictive survival curves for PFS, with corresponding 95% credible intervals, for two new representative patients characterized by similar covariate values and PDGFR samples as in Figure 4, top ($\Delta \geq 0$) and bottom ($\Delta < 0$), respectively.

Table 1

Elicited utilities for intraoperative air leak (IAL) resolution times.

Days to resolve IAL	0	1-5	6-10	11-15	16-20	21-25	26-30	31-35	> 35
Utility	100	50	10	6	5	4	3	2	0

Table 2

Numbers of patients with outcomes of each type following induction therapy, and following salvage therapy.

Induction	Death	Resistant	CR
All patients	69	39	102
FAI	17	17	20
FAI+ATRA	15	13	26
FAI+GCSF	20	4	28
FAI+GCSF+ATRA	17	5	28

Salvage	Death after resistance	Death in CR	Progression after CR	Death after progression
All patients	37	9	93	83
HDAC	25	-	53	-
non HDAC	12	-	40	-