BIOINFORMATICS

Vol. 00 no. 00 xxxx
Pages 1–8

Nonparametric Bayesian Approach to Post-translational Modification Refinement of Predictions from Tandem Mass Spectrometry

Clement Chung[1,2], Andrew Emili[3,4] and Brendan Frey[1,2,3,5*]

[1]Department of Computer Science, [2]Probabilistic and Statistical Inference Group, [3]Banting and Best Department of Medical Research, [4]Donnelly Centre for Cellular and Biomolecular Research, and [5]Department of Electrical and Computer Engineering, University of Toronto, Toronto, Canada.

ABSTRACT

Motivation: Tandem mass spectrometry (MS/MS) is a dominant approach for large-scale high-throughput post-translational modification (PTM) profiling. Although current state-of-the-art blind PTM spectral analysis algorithms can predict thousands of modified peptides (PTM predictions) in an MS/MS experiment, a significant percentage of these predictions has inaccurate modification mass estimates and false modification site assignments. This problem can be addressed by post-processing the PTM predictions with a PTM refinement algorithm. We developed a novel PTM refinement algorithm, *i*PTMClust, which extends a recently introduced a PTM refinement algorithm PTMClust that leverages on the power of a nonparametric Bayesian model that can better account for uncertainties in the quantity and identity of PTMs in the input data. The use of this new modeling approach enables *i*PTMClust to provide a confidence score per modification sites that allows fine-tuning and interpreting resulting PTM predictions.

Results: The primary goal behind *i*PTMClust is to improve the quality of the PTM predictions. First, to demonstrate that *i*PTMClust produces sensible and accurate cluster assignments, we compare it to *k*-means clustering, mixture of Gaussians (MOG) and PTMClust on a synthetically generated PTM dataset. Second, in two separate benchmark experiments using PTM data taken from a phosphopeptide and a yeast proteome study, respectively, we show that *i*PTMClust outperforms state-of-the-art PTM prediction and refinement algorithms, including PTMClust. Finally, we illustrate the general applicability of our new approach on a set of human chromatin protein complex data, where we are able to identify putative novel modified peptides and modification sites that may be involved in the formation and regulation of protein complexes. Accurate PTM profiling is an important step in understanding the mechanisms behind many biological processes and should be an integral part of any proteomic study.

Availability: Our algorithm is implemented in Java and is freely available for academic use from http://www.psi.toronto.edu/ cchung/iPTMClust/

Supplementary Information: Supplementary data are available at Bioinformatics online or at http://www.psi.toronto.edu/iPTMClust/.

Contact: frey@psi.utoronto.ca

1 INTRODUCTION

Post-translational modifications (PTMs) are known to play a vital role in the cell, and are proven to be instrumental in many disease-related studies (Lehninger *et al.*, 1993). A core task in studies involving PTMs is PTM prediction, i.e., identification of peptide sequences and PTMs associated with each modified peptide within a biological sample. A preferred experimental procedure for PTM prediction is MS/MS followed by an analysis with a blind PTM search engine. Reviews of protein mass spectrometry and the detection of PTM by mass spectrometry can be found in (Domon and Aebersold, 2006; Cantin and Yates, 2004). Blind PTM search engines are commonly used because of their versatility to account for both known and novel PTMs. However, PTM predictions produced by blind PTM search engines alone is insufficient since a significant percentage of these PTM predictions contains inaccurate modification masses and incorrect modification positions (Keller *et al.*, 2002; Ramakrishnan *et al.*, 2009a,b). The fragmentation process is often incomplete, and the presence of labile PTMs may interfere with this process (Mikesh *et al.*, 2006). Both issues combined result in spectra missing peaks that in turn may lead to ambiguous or erroneous modification predictions. The presences of natural stable isotopes, such as carbon-13, in addition to electronic noise are major contributors to inaccurate mass measurements. These issues are more prominent in spectra generated from low mass resolution mass spectrometers (e.g., ion trap mass spectrometers), which are still commonly used in today's MS studies. Therefore, it is prudent to incorporate PTM refinement as part of a PTM prediction pipeline, as it can significantly improve the quality of PTM predictions. Previous studies demonstrate that post-processing greatly improves the number of positive predictions while reducing the amount of false PTM assignments (Chung *et al.*, 2011; Tanner *et al.*, 2008).

PTM refinement can be classified into two types of approaches, one that scores the localization of PTMs and one that refines observed modification masses and modification positions. The first type provides a way to evaluate the quality of predicted modification sites from PTM search engines. The two main strategies for scoring the reliability of modification site localizations are: 1) to calculate the probability that a peak responsible for the site determination is matched at random and 2) to compute the search engine score difference between predictions with varying site localizations. Methods that use the former strategy include A-score (Beausoleil *et al.*, 2006), PTM Score (embedded in MaxQuant and Andromeda) (Olsen *et al.*, 2006), the Phosphorylation Localization Score (PLS) in InsPecT (Albuquerque *et al.*, 2008), SLoMo (Bailey *et al.*, 2009), Phosphinator (Phanstiel *et al.*, 2011), PhosphoRS (Taus *et al.*, 2011). Examples of the latter scoring strategy are Mascot Delta Score (Savitski *et al.*, 2011), the SLIP score in Protein Prospector (Baker *et al.*, 2011) and the variable modification locallization (VML) score in Spectrum Mill (Agilent, 2005). A review of the different modification site

*to whom correspondence should be addressed

scoring localization methods is provided in (Chalkley and Clauser, 2012).

PTM refinement using a modification site localization scoring algorithm can be achieved by reassigning the modification position to the highest scoring position for each modified peptide. However, modification site localization scoring methods are limited when used for general PTM refinement due to the following three reasons. First, a predefined list of PTMs is required for these scoring methods. Second, these scoring methods assume that input predicted modification masses are error-free and are mapped precisely to one of the PTMs in the predefined list. Lastly, most of these scoring methods are designed to score only phosphorylated predictions. Consequently, modification site localization scoring methods are ill-suited to analyze PTM datasets generated from blind PTM search engines.

The second type of PTM refinement approaches is to refine both observed modification masses and modification positions. Two recently published algorithms that use this type of PTM refinement method are PTMFinder (Tanner *et al.*, 2008) and our previous algorithm, PTMClust (Chung *et al.*, 2011). PTMFinder takes a peptide level approach to PTM refinement, where it groups and reanalyzes spectra mapping to the same modified peptide sequence to produce for each spectrum a final peptide sequence with a modification mass and a modification position. Hence, refinement using this method is limited to modified peptides that occur multiple times in the same dataset. As shown in a recent study, a modified peptide is rarely found more than three times even for a large-scale, genome-wide experiment (Chung *et al.*, 2011). Furthermore, PTMFinder suffers from favouring high abundance modified peptides and discretizing observed modification masses. Alternatively, PTMClust accounts for the errors of the observed modification masses and modification positions at the PTM level that overcomes many of the issues with PTMFinder.

The principle behind and the distinguishing feature of PTMClust is modelling modifications at the PTM level instead of at the peptide level. This approach has the advantage of allowing the model to account for low abundance modified peptides since other peptides with the same underlying PTM can help identify the correct but unknown modification mass and modified amino acid. PTMClust uses a generative model to capture the hidden relationship between factors influencing the PTM mapping process. PTMClust uses the EM algorithm and a modified version of the split and merge model selection method to learn and infer an optimal parameter setting for the model. As part of the model selection procedure, a range of models are learned by adjusting a model complexity parameter, and the final model is selected by weighting the trade-off between false positives (decoy peptides) allowed and real peptides detected. Despite this cumbersome procedure, the resulting PTM predictions are of higher quality than those taken from existing blind PTM search engines alone or post-analyzed with PTMFinder (Chung *et al.*, 2011).

Although it produces class-leading results, PTMClust has its limitations. In our new algorithm, infinite PTMClust (*i*PTMClust), we set out to address three specific drawbacks of PTMClust: 1) the use of a greedy-based, non-automatic model selection algorithm, 2) the need for manual parameter tuning and 3) the lack of a confidence score per modification position. We overcame these issues by extending the PTMClust model to allow for an unbounded number of mixture components that can account for uncertainties

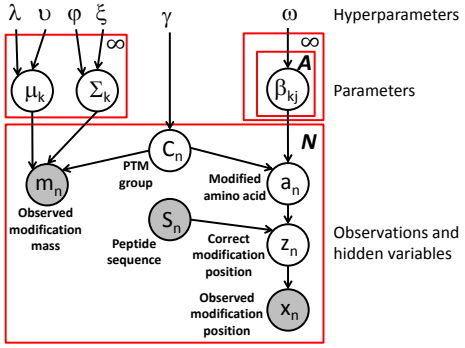

Fig. 1. A Bayesian network describing the generative model for our new algorithm *i*PTMClust, using plate notation. The shaded nodes represent observed variables, the unshaded nodes represent hidden variables and the variables outside the plate are model parameters. The model describes how the observed modification mass and the modification position are generated. The bottom part captures the assumption on how each observation is generated. The plate notation indicates that there are N copies of the model, one for each input data. The top portion outlines the structure of the hierarchy of priors and hyperpriors placed on the model parameters, mixing coefficients, modification mass means, modification variances and probability of modified amino acid. The outer plate shows that there are infinite copies, one for each possible PTM group. The probability of modified amino acid β_{kj} is embedded in two plates signifies that there are $K \times A$ copies, one for each $K \to \infty$ PTM groups and A possible amino acids. The variables outside the plates are hyperpriors' parameters.

in the quantity and identity of PTMs in the input data. This extension parallels the conversion from a finite to an infinite mixture model (IMM), but the complex nature of the underlying PTMClust model makes this extension nontrivial. We derived and implemented the split-merge Metropolis-Hastings (Jain and Neal, 2000) and the Gibbs sampling algorithm (Bishop, 2006; Geman and Geman, 1984) for our model to efficiently infer the groupings of input modified peptides and refine the peptides' modification masses and modification positions. At the end, *i*PTMClust achieves the following benefits: 1) outperforming PTMClust and other PTM refinement algorithms, 2) providing a fully-automated model selection method without the need for any manual parameterization and 3) offering modification position level confidence scores that users can use to assess the quality of the result and further refine their analyses. In a series of benchmark experiments on both synthetic and real-world phosphopeptide datasets, we show that *i*PTMClust outperforms PTMClust and other state-of-the-art PTM prediction and PTM refinement algorithms. To ensure broad applicability, we have designed and optimized *i*PTMClust to analyze PTM data generated from both low and high resolution MS/MS spectra processed by popular blind PTM search engines. Same as in PTMClust, the input to *i*PTMClust is a list of PTM predictions consisting of the peptide sequence, modification position and modification mass.

2 METHODS

Similar to PTMClust, at the heart of *i*PTMClust is a generative probability model that describes a process in which observed modified peptides can be generated by modelling the complex interactions between hidden variables that play a role in the protein modification process. Given an observed modification, we assume it comes from one of many PTMs. However, the number and identity of these PTMs are unknown. Our method accounts for this uncertainty by considering as many different PTMs as needed, represented by an infinite number of PTM groups. By defining appropriate priors on the hidden variables, over-fitting can be avoided, and only a finite

set of PTM groups are used at anytime during inference. The latter point is important because it makes calculations in the algorithm tractable. During inference, the properties of the active PTM groups are influenced by the input data and the chosen priors. We adapted both the Gibbs sampling and the restricted Gibbs sampling split-merge algorithms to infer the values of the hidden variables and parameters in our model. After inference, these hidden variables and model parameters can be used to deduce the true modification mass and a confidence score per possible modification position for each input peptide sequence, or the maximum *a posteriori* (MAP) estimate of modification masses and positions.

2.1 *i*PTMClust Algorithm

The core of the generative model in *i*PTMClust is the same as in PTMClust: it describes how a pair of observed modification mass and modification position are generated. *i*PTMClust extends PTMClust by introducing priors on model variables and parameters that govern the choice of active PTM groups from a boundless number of PTM groups. Given the type of PTM (PTM group) chosen from one of the limitless numbers of PTM groups, we can generate the observed modification mass as a noisy version of the modification mass mean, and select an amino acid most likely to contain the modification as the modified amino acid. Given the peptide sequence, we can choose as the 'true' modification position a position uniformly along it that matches the modified amino acid. Finally, we assume that the observed modification position is a noisy version of the 'true' modification position.

The structural relationships between variables are shown by the Bayesian network in Fig. 1. The top part outlines the priors represented by their corresponding hyper-parameters placed on the model parameters: mixing coefficient, modification mass means, modification mass variances and probability of modification occurring on an amino acid. The bottom portion describes the model for one input peptide and is repeated for N inputs, as indicated by the plate notation. The extension to an IMM in *i*PTMClust follows similar steps from a typical mixture model with finite to infinite number of mixing components given in (Escobar and West, 1994; Rasmussen, 2000). In the following, we outline the rational and intuition behind the different components of our model, and the conditional probability equations of model variables used in the derived inference methods. Detail derivation can be found in the Appendix.

In our model, each input peptide sequence S_n, indexed by $n \in \{1, ..., N\}$, where N is the number of peptides in the dataset, has a corresponding discrete peptide length L_n, observed modification position $x_n \in \{1, ..., L_n\}$, and observed modification mass m_n. We denote the amino acid in position j of the input sequence n as $S_n(j)$. The total number of values $S_n(j)$ can take on is $A = 24$, which includes the 20 naturally-occurring amino acids and four special characters indicating the beginning and end of proteins and peptides. Additionally, we denote o_k to be the number of input peptides assigned to cluster k. The hidden variable $c_n \in [1, .., \infty]$ denotes the unknown PTM group that peptide sequence n is assigned to. Following from the derivation of IMMs with a Chinese Restaurant Process (CRP) (Ferguson, 1973; Antoniak, 1974), its probability conditioned on all other c_n's takes into consideration the likelihood of the n-th input modified peptide belonging to the PTM group index is given as

$$P(c_n = k | c_{N\backslash n}, \gamma, \Theta) = \begin{cases} b\frac{o_{-n,k}}{N-1+\gamma} & \text{if } o_{-n,k} > 0, \\ b\frac{\gamma}{N-1+\gamma} & \text{otherwise,} \end{cases} \quad (1)$$

where $o_{-n,k}$ is the number of peptides assigned to cluster k that does not consider n-th peptide sequence, $N\backslash n$ indicates all indices excluding n, $C_{N\backslash n}$ is shorthand notation for $C_i : \forall i \in \{N\backslash n\}$, b is the appropriate normalizing constant so the probabilities sum to one, γ is the hyper-parameter concentration parameter and Θ represents hyper-parameters λ, υ, φ, ξ, γ and ω. Furthermore, we can write the general form of Eq. 1 that takes into consideration the likelihood of the n-th input modified peptide belonging to the PTM group index by c_n as follows:

$$P(c_n = k | c_{N\backslash n}, \gamma, \Theta) \quad (2)$$

$$= \begin{cases} b\frac{o_{-n,k}}{N-1+\gamma} P(a_n, z_n, x_n, m_n | c_n, S_n, \Theta) & \text{if } o_{-n,k} > 0, \\ b\frac{\gamma}{N-1+\gamma} \int P(a_n, z_n, x_n, m_n, \theta | c_n, S_n, \Theta)\, \partial H_0(\theta) & \text{otherwise,} \end{cases}$$

where θ represents parameters μ, Σ and β and H_0 indicates the prior distribution placed on μ, Σ and β. The calculations of the conditional posterior probability $P(a_n, z_n, x_n, m_n | c_n, S_n, \Theta)$ and the integral $\int P(a_n, z_n, x_n, m_n, \theta | c_n, S_n, \Theta)\, \partial H_0(\theta)$ are given later in Eq. 13 and Eq. 14, respectively. Given a vague inverse gamma prior to hyper-parameter γ, $P(\gamma) = \mathcal{IG}(1, 1)$, its conditional posterior can be derived by combining the joint distribution of c_n's with the prior to give

$$P(\gamma | k, N) \propto \frac{\gamma^{k-\frac{3}{2}} \exp\left(\frac{-1}{2\gamma}\right) \Gamma(\gamma)}{\Gamma(N+\gamma)}. \quad (3)$$

We assume that the observed modification mass for each PTM group is noisy and modelled it to be normally distributed around the true modification mass, given as $P(m_n | c_n = k) = \mathcal{N}(\mu_k, \Sigma_k)$, where μ_k and Σ_k are the parameters modification mass means and variances for the k-th PTM group.

For mathematical convenience, we provided the conjugate prior normal-inverse gamma distribution with hyper-parameters mean λ, variance υ, shape φ and scale ξ^{-1} are to all PTM groups. Board and vague corresponding priors are given to each hyper-parameters to account for uncertainty of their values (details are given in the Supplementary). Given combination of prior and hyper-prior distributions used, the conditional posteriors for λ, υ, φ and ξ have the following forms:

$$P(\lambda | \mu_{1\text{-}k}, \upsilon) = \mathcal{N}\left(\frac{\mu_*/\sigma_*^2 + \upsilon \sum_{j=1}^{k} \mu_j}{1/\sigma_*^2 + k\upsilon}, \frac{1}{1/\sigma_*^2 + k\upsilon}\right), \quad (4)$$

$$P(\upsilon | \mu_{1\text{-}k}, \lambda) = \mathcal{G}\left(k+1, \left[\frac{1}{k+1}\left(\sigma_*^2 + \sum_{j=1}^{k}(\mu_j - \lambda)^2\right)\right]\right), \quad (5)$$

$$P(\varphi | \Sigma_{1\text{-}k}, \xi) \propto \Gamma\left(\frac{\varphi}{2}\right)^{\text{-}k} \exp\left(\frac{1}{2\varphi}\right) \left(\frac{\varphi}{2}\right)^{\frac{k\varphi-3}{2}} \prod_{j=1}^{k} (\Sigma_j \xi)^{\frac{\varphi}{2}} \exp\left(\frac{\text{-}\varphi\Sigma_j\xi}{2}\right), \quad (6)$$

$$p(\xi | \Sigma_{1\text{-}k}, \varphi) = \mathcal{G}\left(k\varphi + 1, \left[\frac{1}{k\varphi+1}\left(\sum_{j=1}^{k} \Sigma_j\right)\right]^{-1}\right), \quad (7)$$

where $\Gamma(\cdot)$ is the gamma function with the form $\Gamma(n) = (n-1)!$ for a positive integer n and 1-k is short for $1, ..., k$.

By using conjugate priors on μ_k and Σ_k, both these parameters can be integrated out to give a probability of m_n based directly on the hyper-parameters, given as

$$P(m_n | c_n, \lambda, \upsilon, \varphi, \xi) \quad (8)$$
$$= \int P(m_n | c_n = k, \mu_k, \Sigma_k) P(\mu_k | \lambda, \upsilon) P(\Sigma_k | \varphi, \xi)\, \partial\mu_k \partial\Sigma_k$$
$$= t\left(\hat{\varphi}, \hat{\lambda}, \frac{\hat{\upsilon}+1}{\hat{\upsilon}+o_k}\hat{\xi}\right),$$

where $t(\cdot)$ is the student's t-distribution, o_k is the number of peptides assigned to cluster k, $\hat{\varphi} = \frac{(\varphi\lambda)+(o_k\bar{m}_k)}{\varphi+o_k}$, $\hat{\lambda} = \lambda + o_k$, $\hat{\upsilon} = \upsilon + o_k$, $\hat{\xi} = \xi + \sum_{i:\forall i, c_i=k} (m_i - \bar{m}_k)^2 + \frac{\lambda o_k (\bar{m}_k - \varphi)^2}{\lambda + o_k}$, and $\bar{m}_k = \frac{1}{o_k} \sum_{i:\forall i, c_i=k} m_i$ is the average observed modification mass for peptides assigned to k-th PTM group.

Let $a_n : n \in \{1, ..., N\}$ denote the true (hidden) modified amino acid (i.e., the amino acid that the PTM occurs on) for the n-th peptide sequence. Then, the probability of a_n's given that the PTM group is k is modelled as a multinomial distribution with parameters $\beta_{ki} \forall i = 1, ..., A$. A dirichlet distribution prior, which is the conjugate prior, with hyper-parameter ω is given to the β_{ki}'s. The hyper-parameter ω is given a vague inverse gamma distribution. Since a conjugate prior is used on β_{ki}'s, we integrate out the β_{ki}'s using a standard Dirichlet integral. Thus, it leads to the conditional posterior of a cluster assignment given all others having the form

$$P(a_n = i | c_n = k, a\text{-}_n, z_n, x_n, S_n, \omega) = \frac{o_{\text{-}n,ki} + \frac{\omega}{A}}{o_k - 1 + \omega} P(z_n, x_n | a_n, S_n), \quad (9)$$

where $P(z_n, x_n|a_n, S_n)$ can be factorized into $P(z_n|a_n, S_n)P(x_n|s_n)$. Both $P(z_n|a_n, S_n)$ and $P(x_n|s_n)$ are discussed below. The conditional posterior for the hyper-parameter ω can be derived similar to Eq. 3 to give

$$P(\omega|A, N) = \frac{\omega^{A-3/2} \exp\left(\frac{-1}{2\omega}\Gamma(\omega)\right)}{\Gamma(N+\omega)}. \quad (10)$$

The details of the rest of the model are the same as those in PTMClust given in (Chung *et al.*, 2011). To reiterate, given the peptide sequence S_n and a_n, we modelled each occurrence of a_n in the peptide sequence to have the same probability of being the true (unobserved) modification position z_n, which can be written as

$$P(z_n = j|a_n = i, S_n) = \begin{cases} \frac{1}{\delta_{ni}+1} & \text{if } S_n(j) = i, j > 0, \\ \frac{1}{\delta_{ni}+1} & \text{if } j = 0, \\ 0 & \text{otherwise,} \end{cases} \quad (11)$$

where δ_{ni} denotes the number of times amino acid i occurs in sequence n and $z_n = 0$[1] indicates that the true PTM occurs outside of the given peptide sequence. Given the true modification position z_n, the probability of modification position error $(x_n - z_n)$ for the observed modification position x_n is modelled with a discrete probability distribution, given as

$$P(x_n|z_n = j) = \begin{cases} \phi(x_n - j) & \text{if } j > 0, \\ \phi(L_n) & \text{if } j = 0, \\ 0 & \text{otherwise,} \end{cases} \quad (12)$$

where the likelihood function ϕ accounts for the modification position error. This likelihood function is shared across all PTM groups and was inferred from our empirical observation of the yeast PTM dataset as described in (Chung *et al.*, 2011).

Finally, we outline how the conditional posterior probability $P(a, z, x, m|c, S, \Theta)$ and the integral $\int P(a, z, x, m, \mu, \Sigma, \beta|c, S, \Theta)\, \partial H_0(\mu, \Sigma, \beta)$ used in Eq. 2 can be evaluated analytically. Based on the structure of the Bayesian network of our model given in Fig. 1, for a given PTM group k, $P(a, z, x, m|S, \Theta)$ can be derived as follows:

$$\begin{aligned} &P(a_n, z_n, x_n, m_n|c_n, S_n, \Theta) \quad (13) \\ &= \int P(a_n, z_n, x_n, m_n, \mu, \Sigma, \beta|c_n, S_n, \Theta)\, \partial H_{-n,k}(\mu, \Sigma, \beta) \\ &= P(z_n|a_n, S_n)P(x_n|z_n)\Big[Dir(\omega + o_k)\Big]\left[t\left(\hat{\varphi}, \hat{\lambda}, (\hat{\upsilon}+1)\hat{\xi}/(\hat{\upsilon}+o_k)\right)\right], \end{aligned}$$

where $H_{-n,k}$ is the posterior distribution of μ, Σ and β based on their priors and all peptides that are assigned to the k-th PTM group excluding the n-th peptide sequence, and the variables $\hat{\varphi}$, $\hat{\lambda}$, $\hat{\upsilon}$ and $\hat{\xi}$ are defined before in Eq. 8. It is easy to see the integral with respect to priors $\int P(a_n, z_n, x_n, m_n, \mu, \Sigma, \beta|c_n, S_n, \Theta)\, \partial H_0(\mu, \Sigma, \beta)$ for unoccupied clusters is equivalent to setting $o_k = 0$ in Eq. 13, which can be written as

$$\begin{aligned} &\int P(a_n, z_n, x_n, m_n, \mu, \Sigma, \beta|c_n, S_n, \Theta)\, \partial H_0(\mu, \Sigma, \beta) \quad (14) \\ &\quad = P(z_n|a_n, S_n)P(x_n|z_n)\Big[Dir(\omega)\Big]\left[t\left(\hat{\varphi}, \hat{\lambda}, (\hat{\upsilon}+1)\hat{\xi}/\hat{\upsilon}\right)\right]. \end{aligned}$$

By combining the structure of the Bayesian network and the conditional distributions described above, we can write the joint distribution as

$$\begin{aligned} &P(c, a, z, x, m, \mu, \Sigma, \beta, \lambda, \upsilon, \varphi, \xi, \gamma, \omega|S, \Psi) \quad (15) \\ &= \prod_{n=1}^{N} \Big[P(\gamma|\Psi)P(c_n|\gamma)P(\lambda|\Psi)P(\upsilon|\Psi)P(\varphi|\Psi)P(\xi|\Psi)P(\omega|\Psi) \\ &\quad P(m_n|c_n, \lambda, \upsilon, \varphi, \xi)P(a_n|c_n, \omega)P(z_n|a_n, S_n, \Psi)P(x_n|z_n, \Psi)\Big], \end{aligned}$$

where Ψ represents the model hyper-parameters for the hyperpriors placed on γ, λ, υ, φ, ξ, and ω.

[1] $z_n = 0$ is needed to avoid numerical issues since our algorithm considers each amino acid as a possible modification target.

2.2 Inference Method

The combination of complicated interactions of hidden variables and priors in *i*PTMClust leads to a complex joint distribution over high-dimensional spaces, which is impossible to characterize analytically in its entirety. This prevents the application of optimization-based inference methods like the EM algorithm. Instead, we used a Markov Chain Monte Carlo (MCMC) approximate inference method (Escobar and West, 1994; Neal, 2000b). MCMC methods are commonly used for IMM models. In MCMC sampling, the model posterior distribution can be cleverly sampled to collect instances of parameter and variable settings likely under the model. Given a large enough collection of samples, the posterior can be approximated, and the ideal settings of model parameters and hidden variables can be inferred.

Although the Gibbs sampling algorithm is a commonly used MCMC sampling method for nonparametric Bayesian clustering models such as ours, it can mix poorly and produce poor results when the input dataset is large and complex (Jain and Neal, 2000). Hence, in addition to the Gibbs sampling method, we derived and implemented the restricted Gibbs sampling split-merge algorithm (or the split-merge sampling algorithm for short) (Jain and Neal, 2000) for *i*PTMClust. The split-merge sampling algorithm is designed to mitigate the aforementioned issues with the Gibbs sampling method. We adhered to the recommended settings for running this sampling algorithm given in (Jain and Neal, 2000).

Both the Gibbs sampling split-merge algorithm and Gibbs sampling methods require to perform Gibbs sampling steps. For each Gibbs sampling step, we employ the following procedure. First, we sample the set of hidden variables (parameters are considered as hidden variables) associated with an input peptide n, where $n \in [1, ..., N]$. Starting with the n-th input observation, for each represented (occupied) cluster K_{rep}, we draw a_n according to Eq. 9 and z_n using Eq. 11. Next, we sample a new cluster assignment c_n from Eq. 2. Finally, we update the sufficient statistics associated with assignments for c_n and a_n. We repeat this procedure for each input peptide sequence. At the end of each Gibbs sampling step we obtain a new value for each hyper-parameter, with the exception for γ, ω and ξ, directly by sampling from their conditional posterior distribution given in Eq. 4, Eq. 5 and Eq. 6, which are all distributions of standard form. Although they are not standard distributions, the conditional posteriors for γ, ω and ξ are all log-concave, which implies that they are unimodal in the logarithmic domain and have a single global optimum. The log-posteriors of γ, ω and ξ are given as

$$\ln P(\gamma|k, N) = C + \left(k - \frac{3}{2}\right)\ln(\gamma) - \frac{1}{2\gamma}\ln\Gamma(\gamma) - \ln\Gamma(N+\gamma), \quad (16)$$

$$\ln P(\omega|A, N) = C + \left(k - \frac{3}{2}\right)\ln(\omega) - \frac{1}{2\omega}\ln\Gamma(\omega) - \ln\Gamma(N+\omega), \quad (17)$$

$$\begin{aligned} \ln P(\xi|\Sigma_1, ..., \Sigma_k, \varphi) &= C - k\ln\Gamma\left(\frac{\xi}{2}\right) - \frac{1}{2\xi} + \quad (18) \\ &\frac{k\xi - 3}{2}\ln\frac{\xi}{2} + \sum_{k=1}^{K_{rep}} \frac{\xi}{2}(\ln\Sigma_k + \ln\varphi) - \frac{\xi\Sigma_k\varphi}{2}, \end{aligned}$$

where C is a normalizing constant. Given the equations for the log-posteriors and their log-concave property, a new value for each hyper-parameter can be efficiently sampled using the slice sampling method (Neal, 2000a), which is an efficient sampling method and is simple to implement.

Ideally, an unlimited number of MCMC samples should be collected to fully and accurately estimate the posterior. However, in practice, it is common to terminate the sampling procedure after a fixed number of sampling iterations at the cost of posterior estimation accuracy. The number of finite samples to collect is based on an analysis of how well the samples have mixed by evaluating trace plots of the distribution of posterior probabilities and the number of clusters over time. Furthermore, common to all MCMC methods is a burn-in period at the beginning during which the sampling algorithm mixes poorly. The samples in this burn-in period need to be removed. The number of burn-in sampling iterations to remove from consideration is also determined by examining the trace plots. Despite being data dependent, we found that 1,000 burn-in samples and a total of 15,000 samples are enough to produce a good approximation of the posterior for

large-scale PTM datasets, such as the phosphopeptides, yeast proteome and human protein-protein interaction datasets that we used in this study. For simpler data, such as the synthetic and phosphorylation datasets, a setting of 100 burn-ins and 6,000 total samples is sufficient. A detailed analysis is provided in the Supplementary. Lastly, to counter auto-correlations amongst the samples, we only use results taken from every fifth sample.

2.3 Background Model

Unlike PTMClust, the background model for *i*PTMClust does not explicitly encompass a predefined background component. Instead, it uses a background model consisting of multiple background components learned directly from the input data. Based on the modification mass variance calculated for each PTM group, we define a background component to be a PTM group with a variance $\geq$ 2.0. This threshold is chosen based on the assumption that the variance for each PTM group is small, because we believe that for a PTM group to be physically relevant, it should have a well-defined modification mass. Therefore, PTM groups with a large modification mass variance are believed to contain spurious data. By allowing for multiple background components instead of one, we observe empirically that the new model is better at capturing spurious data.

3 RESULTS

Our first goal is to demonstrate that *i*PTMClust outperforms existing algorithms both in terms of finding correct clustering assignments and in refining PTMs on a set of noisy modified peptide sequences. To this end, we conducted two experiments. First, a benchmark of *i*PTMClust versus standard clustering algorithms, *k*-means and MOG, as well as our foregoing algorithm PTMClust that is discussed in detail in the Supplementary. The key observation, given in Supplementary figure S1, is that *i*PTMClust using the split-merge algorithm attains the most consistent results across different settings and achieves increasingly better results than PTMClust as the problem becomes more complex. Given the true modification positions are known, we can evaluate further how well PTMClust and implementations of our new algorithm perform for the task of PTM refinement. Supplementary figure S2 shows the same trend as above: *i*PTMClust with the split-merge algorithm generally outperforms the others in almost all cases. The second experiments conducted is a comparison of blind PTM search engines SIMS, InsPecT and MODmap, a state-of-the-art PTM refinement algorithm PTMFinder and our algorithms PTMClust and *i*PTMClust on detecting the true modification positions from a well-studied phosphopeptide dataset. In each of the experiments, we utilize both the split-merge and the Gibbs sampling inference algorithms for *i*PTMClust. We report results based on MAP estimation for *i*PTMClust for both experiments. In the second experiment, we explore the advantage of interpreting results with an averaging over samples approach.

The second goal is to directly show the applicability of *i*PTMClust to datasets taken from studies of complex protein solutions. To achieve this goal, we analyze data taken from a genome-wide yeast and a human chromatin-specific protein complex study. We have limited our analyses to only post-processing PTM predictions generated from SIMS with either PTMClust or *i*PTMClust. We include analysis from PTMClust to highlight that *i*PTMClust is producing sensible results.

In these experiments, we initialize *k*-means, MOG and PTMClust with settings that are outlined in (Chung *et al.*, 2011). The settings for the number of burn-in and total samples for *i*PTMClust are described in the Supplementary.

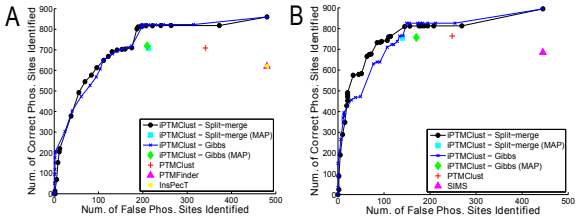

Fig. 2. The figure is a plot of the number of correct phosphorylation sites identified *v.s.* the number of false phosphorylation sites identified for *i*PTMClust and state-of-the-art PTM refinement and prediction algorithms a phosphopeptide dataset with known modification sites. (a) compares *i*PTMClust against PTMClust, InsPecT and PTMFinder. The PTM predictions output from InsPecT are used as the baseline and post-processed by *i*PTMClust, PTMClust and PTMFinder. (b) compares *i*PTMClust against PTMClust and SIMS with output from SIMS are used as the baseline. The curves for *i*PTMClust using the split-merge and the Gibbs sampling algorithm are produced by calculating a confidence score per modification position through inference by averaging over samples and varying the confidence score threshold [0,1]. These methods for evaluating the result outperform their counterparts using MAP estimation. More importantly, *i*PTMClust using any one of the inference methods achieves better results than InsPecT, SIMS, PTMClust and PTMFinder.

3.1 Benchmarking Against Phosphopeptide Predictions

We compare *i*PTMClust against current PTM search engines and PTM refinement algorithms using a real-world dataset enriched for phosphopeptides containing modification sites that are validated (Beausoleil *et al.*, 2004). We will refer to the identities of the known peptide sequences and their modification sites as the reference. The dataset consists of 1655 spectra, but we will focus exclusively on the 1340 spectra mapped and curated as singly modified phosphopeptides (SIMS, InsPecT, PTMClust and *i*PTMClust are limited to one modification per peptide sequence). In this analysis, we define positives (P) as outputs from the base blind PTM search engine that match to the reference considering only their peptide sequence, i.e., disregarding the positions of their modification, and negatives (N) as all other outputs that do not match their corresponding reference peptide sequences. Each blind PTM search engine produces a different number of P and N. For SIMS, PTMClust and *i*PTMClust (SIMS was used as the base blind PTM search engine), there are 895 P and 445 N. Lastly, for PTMFinder, which uses InsPecT as its base unrestricted PTM search engine, there are 860 P and 480 N.

For *i*PTMClust, in addition to the MAP estimate, we considered inference by averaging over samples to produce a confidence score per modification position for each output peptide sequence. By varying the confidence threshold setting, we can adjust the sensitivity and specificity of the PTM predictions. For *i*PTMClust results evaluated using confidence scores, we define a PTM prediction as a peptide sequence and its modification positions with confidence scores above the threshold; peptide sequences that do not have at least one modification position with a confidence score above the threshold are considered to be assigned to the background model. For PTMClust and *i*PTMClust using MAP estimate, a prediction is any peptide sequence not assigned to the background model. All peptide sequences assigned to the background model are removed from the evaluation. Lastly, we consider all outputs as predictions for the other algorithms since they do not employ a background model.

Given the true phosphorylation of each peptide are known, we plot the number of correct phosphorylation sites versus the number of incorrect phosphorylation sites identified in Fig. 2a for the results of *i*PTMClust, PTMClust, PTMFinder and InsPecT. For this experiment, the blind PTM search engine InsPecT was used to analyze the input spectra, then each of the PTM refinement algorithms, *i*PTMClust, PTMClust and PTMFinder, were used to post-process the PTM predictions output from InsPecT. For *i*PTMClust, we included results for both the split-merge and Gibbs sampling methods using either MAP estimation or inference by averaging over samples. We varied the confidence threshold to obtain a series of results for *i*PTMClust using inference by averaging over samples. Four key observations can be made from the result: 1) *i*PTMClust outperforms all other algorithms, including PTMClust and PTMFinder; 2) *i*PTMClust with inference by averaging over samples produces better results than its MAP estimate counterpart; 3) a background model is essential for collecting spurious modifications; and 4) *i*PTMClust with the split-merge inference algorithm performs similarly to *i*PTMClust with the Gibbs sampling algorithm except at the region with low number of incorrect phosphorylation sites identified, where the Gibbs sampling method is performing marginally better than its counterpart. The first observation reinforces our conclusion from the study on synthetic data that *i*PTMClust outperforms other PTM refinement algorithms. The second highlights the advantage of providing a confidence score per modification position, where adjusting for the confidence threshold allows us to achieve improved results over MAP estimation. Finally, the results highlight a major benefit of *i*PTMClust beyond PTM refinement: i.e., removal of noisy data.

Next, we conducted the same experiment by post-processing the output from SIMS with PTMClust and *i*PTMClust to ensure that our algorithms are not bias towards any one particular blind PTM search engine. Similarly, we plot the number of correct phosphorylation sites versus the number of incorrect phosphorylation sites identified in Fig. 2b for the results of *i*PTMClust, PTMClust and SIMS. Since PTMFinder is tightly integrated into the InsPecT algorithm, we were not able to decouple PTMFinder from InsPecT to post-process the output from SIMS. Therefore, PTMFinder is omitted from this analysis. The results based on the PTM predictions generated by both SIMS and InsPecT (discussed above) show the same three trends: 1) *i*PTMClust outperforms all other algorithms; 2) *i*PTMClust with inference by averaging over samples produces better results than its MAP estimate counterpart; and 3) a background model is essential for collecting spurious modifications. Specific to this experiment based on the output from SIMS, we observe that *i*PTMClust with the split-merge inference algorithm markedly performs better than *i*PTMClust with the Gibbs sampling in two conditions: when MAP estimate is used and at the region between 30 to 145 number of false phosphorylation sites identified when using inference by averaging over samples.

3.2 Large-scale PTM Analysis of Yeast Proteome

Through a series of benchmark experiments, we have shown that *i*PTMClust beats state-of-the-art algorithms, including our own PTMClust. Furthermore, we demonstrate that *i*PTMClust using the split-merge sampling method produces improved results over the one using the Gibbs sampling method. Next, we will test *i*PTMClust's versatility in detecting diverse PTM groups by applying it to analyze a large-scale PTM dataset taken from analyses of the yeast proteome (LC-MS/MS spectra only) (Krogan *et al.*, 2006) using SIMS. Briefly, the yeast dataset consists of over 2 million ion trap MS/MS spectra of which 19,560 putatively modified peptides were identified by SIMS with modification range (0, 200] Da. The estimated false discovery rate for the predictions made by SIMS is 4.3% based on the number of decoy peptides identified. Here, we present the result taken from our post-analysis using *i*PTMClust with the split-merge algorithm (MAP estimate was used). MAP estimate is used to simplify the analysis and the mapping to known PTMs. The specific setting used for PTMClust is described in (Chung *et al.*, 2011).

Table 1. Summary of peptides with known PTM sites in the yeast proteome dataset.

PTM	Peptides with Known PTM Sites (Num. of Unique Sites)			
	i PTMClust Split-merge	*i* PTMClust Gibbs	PTMClust	SIMS
Phosphorylation	196 (103)	185 (101)	115 (66)	100 (61)
Acetylation	59 (5)	78 (9)	75 (9)	72 (8)
Cysteine Oxidation (Cysteine sulfinic acid)	7 (2)	6 (1)	7 (1)	6 (1)
Others	24 (2)	35 (5)	35 (5)	35 (5)
Total	286 (112)	304 (116)	232 (81)	213 (75)

The known set of modifications was taken from Uniprot (Release 2010_11). We matched the sets of modified peptides produced by SIMS and post-processed with *i*PTMClust with the split-merge algorithm, *i*PTMClust with the Gibbs sampling algorithm and PTMClust to the set of known yeast modification sites. This table shows the number of peptides and unique sites that are mapped to known PTMs. The results show *i*PTMClust is able to improve upon SIMS and significantly outperform PTMClust in a complex dataset.

A summary of commonly known PTMs taken from the Uniprot knowledgebase (Release 2010_11) that are found in our dataset is shown in Table 1. Overall, *i*PTMClust using either the split-merge or the Gibbs sampling method is able to reposition a large portion of modifications to known PTM sites that were missed by SIMS originally (increase of $\sim$ 49% unique PTMs and $\sim$ 34% modified peptides for split-merge, and $\sim$ 55% and $\sim$ 43% for Gibbs). This represents a significant increase over what can be achieved using PTMClust. The most improvement is gained with phosphorylation sites, where post-analysis with *i*PTMClust is able to identify $>$ 65% more (known) unique sites and almost double the number of instances of phosphopeptide when compared to the result obtained from SIMS. However, *i*PTMClust using the split-merge method incorrectly places a number of peptides with acetylation and other modifications in the background model that SIMS correctly identified. We note that *i*PTMClust using the Gibbs sampling method and PTMClust did not make this mistake. A closer look reveals that many of these instances belong to a few unique peptide sequences. The analysis on the yeast proteome dataset confirms that *i*PTMClust can detect other PTMs such as acetylation and cysteine oxidation (cysteine sulfinic acid) in addition to phosphorylation. Moreover, the results reiterate that *i*PTMClust using either the split-merge or the Gibbs sampling method can refine a greater number of PTMs than PTMClust.

3.3 Analysis of Human Protein-protein Interaction Data

Protein complexes and protein-protein interactions studies are a major focal point in the field of proteomics. However, to date,

the focus has been mainly on finding complex memberships and interaction partners. Since it is well established that PTMs, such as phosphorylation and acetylation, play a vital role in the formation and regulation of protein-protein interactions, we seek to complement these studies with an emphasis on the identification of PTMs.

The dataset we used consists of high mass-resolution MS/MS spectra (Orbitrap mass spectrometer) from a human protein-protein interaction study searched using SIMS. The study is a collaboration with the Emili lab at the University of Toronto, and the dataset is not yet published at the time of writing. The experimental protocol used is tandem affinity purification (TAP) (Rigaut *et al.*, 1999; Puig *et al.*, 2001) followed by MS approach. Briefly, the method proceeds by placing a biological tag on all instances of a member of the complex of interest. Next, these tagged proteins are isolated and purified along with their interacting proteins, and finally, the set of purified proteins are subjected to MS analyses.

Here, we chose to focus on three well-studied protein complexes, the Mediator (MED), the RNA Polymerase II (POL2) and the Polycomb Repressive Complex 1 (PRC1). This dataset consists of 17 experiments. These experiments include technical replicates, covering six of 12 proteins known to be in the MED complex, two experiments covering two of the 32 members of the POL2 complex and three experiments covering two of 12 members of the PRC1 complex. There is a total of 13,221 spectra mapped to modified peptides by SIMS (estimated false discovery rate of 13.5% based on the number of decoy peptides identified) with a modification range (0, 300] Da.

Supplementary table T1 highlights the 48 distinct peptide sequences putatively identified to be modified either with acetylation or phosphorylation. Including duplicates, a total of 114 modified peptides was found. Of the 48 unique peptides, 9 had at least one instance with their modification site corrected by *i*PTMClust using the split-merge sampling algorithm that was originally misplaced by SIMS. These are highlighted in bold. We speculate this smaller number improvement is due to two major reasons: 1) a majority of the listed PTMs are acetylation, which SIMS does a good job with, as seen in the yeast study above, and 2) spectra are cleaner due to the use of a high-resolution mass spectrometer. In this list, there are 15 putative novel modification sites and 33 known ones according to the Uniprot knowledgebase (Release 2011_12). Although the complete list is not shown, after removing those assigned to the background model, we have identified a total of 10,409 putative modified peptides such that a large portion of them is mapped to regions in their respective proteins that do not contain known PTMs. Similar to those listed putative, novel phosphorylated and acetylated peptides, we believe this list also contains many high-quality, new PTM discoveries. Hence, this list represents a filtered list of high-quality candidates for further investigation. We have shown a PTM prediction pipeline comprises of a blind PTM search engine and *i*PTMClust can be fruitful in novel discoveries and should be use routinely.

The use of a high mass-resolution mass spectrometer is expected to reduce errors and potentially remove the need for refinement of measured modification masses. Even so, we noticed many instances where the observed modification mass deviates from our refined modification mass by $\sim$1 Da. This modification mass error is believed to be due to the presence of isotopes. Although heuristics can be used to account for these mass shifts, such methods can be error-prone and cannot adapt to unforeseen mass errors. Our results show that *i*PTMClust can handle such errors and improve the quality of PTM predictions taken from an analysis of a high mass-resolution mass spectrometer.

4 CONCLUSION

Accurately identifying PTMs and their potential roles in clinical studies such as biomarker discovery and drug development is an important task. Although thousands of PTM candidates have been reported using blind PTM search engines (Liu *et al.*, 2008; Tanner *et al.*, 2005; Tsur *et al.*, 2005; Liu *et al.*, 2006; Searle *et al.*, 2006; Han *et al.*, 2005), these blind PTM search algorithms suffer from mass measurement inaccuracy and uncertainty in predicting modification positions, making the findings error prone. The importance of post-processing PTM predictions using a PTM refinement algorithm have been established in (Chung *et al.*, 2011; Tanner *et al.*, 2008). The previous state-of-the-art algorithm, PTMClust, achieves a significantly higher PTM prediction accuracy over blind PTM search engines alone and outperforms existing PTM refinement algorithm, PTMFinder. Despite significant improvements in PTM prediction, PTMClust has three main drawbacks of particular interest: 1) it employs of a greedy-based, non-automatic model selection algorithm, 2) it requires manual parameterization on the maximum number of PTM groups and 3) it does not provide a confidence score per modification position.

To address these issues, we introduce *i*PTMClust. *i*PTMClust extends PTMClust by using an infinite mixture model approach that achieves the following three benefits: 1) outperforming PTMClust and other PTM refinement algorithms, 2) providing a fully-automated model selection method without the need for any manual parameterization and 3) offering modification position level confidence scores that users can use to assess the quality of the results and to greater refine their analyses. Through a series of benchmark experiments using both synthetic and real (phosphopeptides and yeast proteome) data, we demonstrated that *i*PTMClust better models the PTM generative process and outperforms PTMClust, PTMFinder and other blind PTM search engines. In addition, we analyzed data generated from a yeast proteome study using *i*PTMClust in which we reported an improvement over the base blind PTM search algorithm SIMS in detecting annotated PTMs. Thousands of putative PTMs were found in this analysis. Moreover, in our in-depth look at PTM predictions for three human protein complexes, MED, POL2 and PRC1, *i*PTMClust identified numerous validated and putative phosphorylated and acetylated peptides that may be involved in the formation and regulation of protein-protein interactions. Further investigations are warranted, but we believe a number of these putative predictions are valid PTMs and can serve to further our understanding of the complexities involved in protein-protein interactions. To summarize, our new algorithm *i*PTMClust is easy to use, achieves overall greater performance than the state of the art, provides confidence scores at the modification position level that allow for a higher flexibility when evaluating potential PTMs and is designed to be broadly applicable to PTM predictions generated from any blind PTM search engine.

Given the rapid advancement of mass spectrometer technology, how applicable is *i*PTMClust going forward? We explore this question by analyzing data generated from high mass-resolution

mass spectrometers, e.g., from an Orbitrap. Although mass accuracy has improved with the use of high mass-resolution mass spectrometers, the presence of isotopes, for example, can result in deviations in observed modification masses. In addition, higher mass accuracy does not necessary equate to errorless modification site determination. We have shown in our analysis of the human protein complex data that problems with mass measurement inaccuracy and uncertainty predicting modification positions continue to exist for data generated from high mass-resolution mass spectrometers, such as an Orbitrap used in the experiment. Our results demonstrate that *i*PTMClust can improve upon PTM predictions taken from data with high mass accuracy, and continue to be a vital component of a genome-wide PTM study.

In designing *i*PTMClust, we have developed two different inference algorithms, the split-merge Metropolis-Hastings and the Gibbs sampling algorithm. While the Gibbs sampling method is the standard inference algorithm for Bayesian mixture models and IMMs, the split-merge method is shown to perform better when dealing with complex datasets (Jain and Neal, 2000). Results of our synthetic data experiments also exhibit this trend. *i*PTMClust using the Gibbs sampling method displays a performance drop while its counterpart using the split-merge sampling algorithm perform consistently well as the data used gets more complex. It is important to note that when there are few PTMs within a small modification mass window, say 1-2 Da (e.g., when two or three PTM groups added in our experiment with the synthetic data), our algorithm using the Gibbs sampling performs at par with or slightly worse than its counterpart using the split-merge method. Furthermore, *i*PTMClust using the Gibbs sampling runs faster per iteration ($\sim 50\%$ quicker) than *i*PTMClust using the split-merge method. Hence, for large datasets, where running time can be overwhelmingly long, we recommend running *i*PTMClust with the Gibbs sampling method as it provides a good trade-off between quality of the result and computational cost.

Despite outperforming its competitors, *i*PTMClust has a number of limitations. Similar to PTMClust, it is unable to handle more than one modification per input peptide sequence and PTM groups identified can contain multiple PTMs if their modification masses are similar. The latter problem is less of an issue when working with high mass-resolution data. Moreover, *i*PTMClust does not consider the underlying spectrum when refining a PTM prediction. The presence of certain peaks in the spectrum can add to support to a residue along the peptide sequence as being the modification position. Its limitations notwithstanding, *i*PTMClust is shown to outperform both PTMClust and previous state of the art in our benchmark tests using both synthetic and real-world PTM data.

ACKNOWLEDGEMENT

We thank Jonathan Olsen, Jian Liu, Yoseph Barash and Leo Lee for valuable conversations and Vincent Fong for his help with collecting the experimental data and generating the SIMS and InsPecT results.

Conflict of Interest: none declared.

REFERENCES

Agilent (2005). Spectrum mill. http://www.chem.agilent.com/en-US/Products-Services/Software-Informatics/Spectrum-Mill/pages/default.aspx.

Albuquerque, C. P. *et al.* (2008). A multidimensional chromatography technology for in-depth phosphoproteome analysis. *Molecular Cell Proteomics*, **7**(7), 1389–1396.

Antoniak, C. E. (1974). Mixtures of dirichlet processes with applications to bayesian nonparametric problems. *The Annals of Statistics*, **2**, 1152–1174.

Bailey, C. M. *et al.* (2009). Slomo: automated site localization of modifications from etd/ecd mass spectra. *Journal of Proteome Res.*, **8**(4), 1965–1971.

Baker, P. R. *et al.* (2011). Modification site localization scoring integrated into a search engine. *Molecular Cell Proteomics*, **10**(7), M111.008078.

Beausoleil, S. *et al.* (2004). Large-scale characterization of hela cell nuclear phosphoproteins. *Proceedings of the National Academy of Science of United States of America*, **101**(33), 12130–12135.

Beausoleil, S. A. *et al.* (2006). A probability-based approach for high-throughput protein phosphorylation analysis and site localization. *Nature Biotechnology*, **24**(10), 1285–1292.

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer.

Cantin, G. T. and Yates, J. R. (2004). Strategies for shotgun identification of post-translational modifications by mass spectrometry. *Journal of Chroma.*, **1053**(1-2), 7–14.

Chalkley, R. J. and Clauser, K. R. (2012). Modification site localization scoring: Strategies and performance. *Molecular Cell Proteomics*.

Chung, C. *et al.* (2011). Computational refinement of post-translational modifications predicted from tandem mass spectrometry. *Bioinformatics*, **27**(6), 797–806.

Domon, B. and Aebersold, R. (2006). Mass spectrometry and protein analysis. *Science*, **312**(5771), 212–217.

Escobar, M. D. and West, M. (1994). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*, **90**, 577–588.

Ferguson, T. S. (1973). A bayesian analysis of some nonparametric problems. *The Annals of Statistics*, **1**(2), 209–230.

Geman, S. and Geman, D. (1984). *Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images*, volume PAMI-6. IEEE.

Han, Y. *et al.* (2005). Spider: software for protein identification from sequence tags with de novo sequencing error. *Journal of bioinformatics and computational biology*, **3**(3), 697–716.

Jain, S. and Neal, R. (2000). A split-merge markov chain monte carlo procedure for the dirichlet process mixture model. *Journal of Computational and Graphical Statistics*, **13**, 158–182.

Keller, A. *et al.* (2002). Empirical statistical model to estimate the accuracy of peptide identifications made by ms/ms and database search. *Analytical Chemistry*, **74**, 5383 – 5392.

Krogan, N. *et al.* (2006). Global landscape of protein complexes in the yeast saccharomyces cerevisiae. *Nature*, **440**(7084), 637–643.

Lehninger, A. L. *et al.* (1993). *Principles of Biochemistry (Second Edition)*. Worth Publishing.

Liu, C. *et al.* (2006). Peptide sequence tag-based blind identification of post-translational modifications with point process model. *Bioninformatics*, **22**(13), e307–e313.

Liu, J. *et al.* (2008). Sequential interval motif search: Unrestricted database searching of global ms/ms datasets for unexpected post-translational modifications. *Analytical Chemistry*, **80**(20), 7846–7854.

Mikesh, L. M. *et al.* (2006). The utility of etd mass spectrometry in proteomic analysis. *Biochim Biophys Acta*, **1764**(12), 1811–1822.

Neal, R. (2000a). Slice sampling. *Annals of Statistics*, **31**, 705–767.

Neal, R. M. (2000b). Markov chain sampling methods for dirichlet process mixture models. *Journal of Computational and Graphical Statistics*, **9**(2), 249–265.

Olsen, J. V. *et al.* (2006). Global, in vivo, and site-specific phosphorylation dynamics in signaling network. *Cell*, **127**(3), 635–648.

Phanstiel, D. H. *et al.* (2011). Proteomic and phosphoproteomic comparison of human es and ips cells. *Molecular Cell Proteomics*, **8**(10), 821–827.

Puig, O. *et al.* (2001). The tandem affinity purification (tap) method: a general procedure of protein complex purification. *Methods*, **24**(3), 218–229.

Ramakrishnan, S. *et al.* (2009a). Mining gene functional networks to improve mass-spectrometry based protein identification. *Bioinformatics*, **25**(22), 2955–2961.

Ramakrishnan, S. R. *et al.* (2009b). Integrating shotgun proteomics and mrna expression data to improve protein identification. *Bioinformatics*, **25**(11), 1397 – 1403.

Rasmussen, C. E. (2000). The infinite gaussian mixture model. In *In Advances in Neural Information Processing Systems 12*, pages 554–560. MIT Press.

Rigaut, G. *et al.* (1999). A generic protein purification method for protein complex characterization and proteome exploration. *Nature Biotechnology*, **17**(10), 1030–1032.

Savitski, M. M. *et al.* (2011). Confident phosphorylation site localization using the mascot delta score. *Molecular Cell Proteomics*, **10**(2), M110.003830.

Searle, B. *et al.* (2006). Identification of protein modifications using ms/ms de novo sequencing and the opensea alignment algorithm. *Journal of proteome research*, **4**(2), 546–554.

Tanner, S. *et al.* (2005). Inspect: identification of posttranslationally modified peptides from tandem mass spectra. *Analytical chemistry*, **77**(14), 4626–4639.

Tanner, S. *et al.* (2008). Accurate annotation of peptide modifications through unrestrictive database search. *Journal of Proteome Research*, **7**, 170–181.

Taus, T. *et al.* (2011). Universal and confident phosphorylation site localization using phosphors. *Journal of Proteome Res.*, **10**(12), 5354–5362.

Tsur, D. *et al.* (2005). Identification of post-translational modifications by blind search of mass spectra. *Nature Biotechnology*, **23**(12), 1562–1567.