Published in final edited form as:

Proceedings (IEEE Int Conf Bioinformatics Biomed). 2016 December ; 2016: 184–189. doi:10.1109/bibm.2016.7822516.

A Multi-Modal Graph-Based Semi-Supervised Pipeline for Predicting Cancer Survival

Hamid Reza Hassanzadeh

Department of Computational Science and Engineering, Georgia Institute of Technology Atlanta, Georgia 30332

John H. Phan, May D. Wang

Department of Biomedical Engineering Georgia Institute of Technology and Emory University, Atlanta, Georgia 30332

Abstract

Cancer survival prediction is an active area of research that can help prevent unnecessary therapies and improve patient's quality of life. Gene expression profiling is being widely used in cancer studies to discover informative biomarkers that aid predict different clinical endpoint prediction. We use multiple modalities of data derived from RNA deep-sequencing (RNA-seq) to predict survival of cancer patients. Despite the wealth of information available in expression profiles of cancer tumors, fulfilling the aforementioned objective remains a big challenge, for the most part, due to the paucity of data samples compared to the high dimension of the expression profiles. As such, analysis of transcriptomic data modalities calls for state-of-the-art big-data analytics techniques that can maximally use all the available data to discover the relevant information hidden within a significant amount of noise. In this paper, we propose a pipeline that predicts cancer patients' survival by exploiting the structure of the input (manifold learning) and by leveraging the unlabeled samples using Laplacian support vector machines, a graph-based semi supervised learning (GSSL) paradigm. We show that under certain circumstances, no single modality per se will result in the best accuracy and by fusing different models together via a stacked generalization strategy, we may boost the accuracy synergistically. We apply our approach to two cancer datasets and present promising results. We maintain that a similar pipeline can be used for predictive tasks where labeled samples are expensive to acquire.

I. INTRODUCTION

Cancer survival prediction provides valuable information for oncologists to choose an effective course of treatment for cancer patients. This is particularly important for terminally ill patients for whom uninformed treatment selection can undermine patients' quality of life [9]. Unfortunately, there is a high amount of stochasticity involved in regard to this clinical endpoint due to numerous confounding factors (e.g., environmental factors) that can play a role in patients' survival. As a result, physicians often make subjective decisions based on their past experience. Yet another major challenge when dealing with clinical and omics data

pertaining to the cancer patients is that a significant portion of these data are unlabeled due to the lack of adequate follow-up information. This is especially the case for the longitudinal studies that span several years (often more than five years,) which results in patient withdrawal from the study or are subject to research budget restrictions. As a result, most research studies filter out these samples and solve the problem within the framework of supervised learning. Even though this strategy works for larger datasets, it fails to produce accurate and robust results in most cancer datasets where small sample sizes remain a bottleneck. This, along with the high dimensionality of the input data, makes the objective difficult to achieve. An alternative to the above-mentioned strategy would be to leverage the information present in the available unlabeled data which are sometimes orders of magnitude bigger in size, compared to the amount of labeled samples. This approach falls into the category of the semi-supervised learning paradigm.

Semi-supervised learning (SSL) is a learning paradigm that takes advantage of both labeled and unlabeled data when training a model. Traditionally, models trained for predictive tasks are either supervised or unsupervised. While the labeled data convey valuable information that can be leveraged to train a model, they are not always available in abundance. On the other hand, the unlabeled data may be readily available in large quantities as they may be cheaper or faster to generate. In the case of cancer data, for example, it is often costly to conduct a long-term research study where the patient eventually dies due to his/her disease. Consequently, the outcome corresponding to the endpoint of study may not be known for a significant part of the dataset. Even though the unlabeled data do not tell us as much about the prediction outcome, may provide some information about the distribution of data that can be exploited to guide model training and lead us to a higher predictive power. However, the key requirement for any SSL method to succeed is that the marginal data distribution has to carry some useful information regarding the inference of the posterior probability [4] (the Bayes rule.) To date, several SSL classification paradigms have been proposed (such as self-training, co-training, transductive SVM [TSVM] and graph-based semi-supervised learning [GSSL]). Despite the success of SSL in different domains, their application to the prediction of clinical endpoints is relatively new. Here, we give an account of the related published studies in chronological order.

As one of the earliest applications of GSSL to tumor classification, Gui et al. [10] used a graph-based approach named local and global consistency (LGC) [27]. LGC is an iterative algorithm inspired from the spreading of activation networks from experimental psychology that can be explained as random walks on graphs. The main idea in LGC is that each unlabeled sample receives some amount of information from its neighbors, which is retained according to some learning rate. In the end, based on the amount of information each unlabeled node receives from other nodes, the label of that node is specified. Note that, one key difference between Laplacian SVM and LGC is that the latter is not considered to be a large margin classifier. Therefore, caution should be practiced when generalization is a concern. Shi et al. [24] used transductive support vector machines (TSVM) to predict recurrence in colorectal cancer patients using microarray gene expression data. TSVM decides the labels for unlabeled samples such that the classification boundary can be placed in a low density region with the margin maximized. This approach, however, is not an appropriate choice when there is no clear boundary between the samples of different classes,

which is often the case for cancer data. A similar approach [19] for prediction of cancer subtypes has also been reported. In another successful application of GSSL, Kim et al. [16] used manifold regularization to enforce smoothness constraints on the intrinsic geometry of the marginal input distribution. They applied their developed method to predict different clinical endpoints such as survival. Moreover, they used different modalities of genomic data and represented each with a graph Laplacian and showed that, by generating a graph Laplacian as a weighted sum of each individual Laplacian, the solution follows the same analytic form. However, due to the quadratic form of the cost function, it could be considered more of a semi-supervised extension to the regularized least squares method and lacks margin generalization. They used a similar approach in their follow-up studies [15], [14]. In another GSSL-related study, Kim et al. [17] used two levels of semi-supervised learning to integrate different layers of omics data. Iteratively, they trained separate GSSL methods on each omics dataset, then used co-training to assign pseudo-labels to the unlabeled samples. More recent published research using similar techniques can be found in [8], [12], [3], [6], [22].

In this article, we develop a pipeline for prediction of cancer survival based on Laplacian support vector machines, an extension of support vector machines to the semi-supervised domain. This method incorporates the geometry of the marginal input distribution and maintains a balance between the loss function and semi-supervised smoothness assumptions where intuitively smoothness means data samples that are closer to each other in the input space, should also be closer in terms of the labels they are assigned. Moreover, we compare and contrast the efficacy of the proposed method as a function of the size of the labeled and unlabeled sets.

II. MANIFOLD REGULARIZATION & LAPLACIAN SVM

Semi-supervised learning differs from supervised learning in that the underlying dataset consists of a set of l labeled inputs $(x_i, y_i)_{i=1}^l$ and a set of u unlabeled samples $(x_i)_{i=l+1}^{l+u}$. As such, there is a probability distribution P on $X \times R$ from which examples are drawn. On the other hand, assuming that the unlabeled samples are drawn from the same underlying stochastic process, they should be distributed according to the marginal distribution P_X of the original P . Graph-based semi-supervised learning (GSSL) [27], on the other hand, is based on the assumption that if the data lie on a manifold of much lower dimensions (the intrinsic geometry of P_X) than that of the input space, then by enforcing smoothness constraints on that manifold for both labeled and unlabeled samples, we can produce a classifier that is more consistent with the input data [5]. GSSL represents each data sample as a node in a weighted graph where the weight attributed to each edge, w_{ij} , is computed according to an affinity function, defined shortly. This results in a dense graph which is often not desired and hence, in a later step, sparsified using techniques such as k -nearest neighbor. The resulting graph serves as a proxy for the manifold and is used within the framework of a variety of manifold regularization techniques. One important matrix that is defined over this graph and which plays an important role in derivation of many of the GSSL based approaches is the graph Laplacian, computed as $L = W - D$ where $W = [w_{ij}]_{i,j}$ is the affinity matrix and the diagonal matrix D is given by $D_{ii} = \sum_{j=1}^{l+u} w_{ij}$ Proposed by Belkin et al. [2],

manifold regularization is a family of semi-supervised learning algorithms that exploits the intrinsic geometry of the marginal distribution P_X by adding an additional regularization term. The support of P_X is assumed to have the geometric structure of a Riemannian manifold M [1]. Within this framework, one can derive an extension to the support vector machine by using the soft margin loss function and thereby incorporating the manifold structure of the input. As a result, the extended classifier not only leverages the manifold structure of the input but also remains a large margin classifier with useful generalization implications that make it an effective tool when dealing with data scarcity. Formally speaking, for any given Mercer kernel [23], $K: X \times X \rightarrow \mathcal{R}$, and the corresponding norm, $\|\cdot\|_K$, a number of popular algorithms including SVM can be cast into the following minimization form with different empirical cost functions:

$$f^* = \arg \min_{f \in H_K} \frac{1}{l} \sum_{i=1}^l V(x_i, y_i, f) + \gamma \|f\|_H^2 \quad (1)$$

where V is some loss function and H_K is the reproducing kernel Hilbert space (RKHS) of functions $X \rightarrow \mathcal{R}$ corresponding to kernel K . The regularization term in equation 1 imposes smoothness conditions on possible solutions. According to the represented theorem the solution to equation 1 exists in H_K and can be expressed as,

$$f^*(x) = \sum_{i=1}^l a_i K(x_i, x) \quad (2)$$

and, as a result, the problem reduces to optimizing over the finite dimensional space of coefficient a_i . The problem with this framework, however, is that it only enforces the smoothness constraint in the kernel space (aka the ambient space) but does not take into account the smoothness of the solution over the manifold, that is, the geometric structure of P_X . To achieve this, they add an extra regularizer to enforce the smoothness of the solution relative to the manifold which expands the minimization problem in equation 1, to the following:

$$f^* = \arg \min_{f \in H_K} \frac{1}{l} \sum_{i=1}^l V(x_i, y_i, f) + \gamma_A \|f\|_H^2 + \gamma_I \|f\|_I^2 \quad (3)$$

where $\|f\|_I^2$ is the norm in the intrinsic space. The regularizer parameters γ_A , γ_I control the smoothness of the solution relative to the ambient and the intrinsic spaces, respectively. It can be shown that when the intrinsic space is a compact manifold $M \in \mathcal{R}^n$ then one natural choice for $\|f\|_I^2$ would be

$$\|f\|_I^2 = \int_M \|\nabla_M f(x)\|^2 dP_x \quad (4)$$

where ∇_M is the gradient of the function over the Riemannian manifold M and it can be approximated with the graph Laplacian corresponding to the proxy graph when the affinity

matrix takes the exponential form $w_{ij} = e(-\frac{\|x_i - x_j\|}{\epsilon}) \forall i \neq j$ and $w_{ii} = 0 \forall i$. This approximation reduces the intrinsic penalty term to

$$\|f\|_I^2 = f^T L f \quad (5)$$

where $f^T = [f(x_1), \dots, f(x_{l+u})]$. Finally, following from the representer theorem [23], they showed that the solution to equation 3 with the soft-margin loss function and equation 5 as the smoothness penalty term on the manifold (henceforth, Laplacian support vector machine) takes the same form as in equation 2 but with a different set of weights α_i .

III. MATERIALS AND METHODS

A. Data

We used RNA-seq data corresponding to cancer tissues for two cancer types, the kidney cancer (KIRC) and the neuroblastoma (NB) which occurs most often in infants and young children. The data for the former cancer was retrieved from the Cancer Genome Atlas (TCGA) data portal (<http://www.tcgadata.nci.nih.gov>) and for the latter we used the data from a previous published study [26]. These data were generated using the Illumina HiSeq 2000 platform. In this study, we use three different feature levels (gene, transcript and junction) derived from these sequencing data. Table I shows the data description for each modality and disease. We divided samples into positive and negative classes according to a selected patient survival threshold. For kidney cancer, this threshold is 5 years and for neuroblastoma cancer it is 9 years. Table II illustrates the statistics of each class and the percentage of unlabeled samples. According to the table, more than half of the samples from each dataset are unlabeled which justifies the use of semi-supervised learning to benefit from them.

B. The proposed pipeline

Figure 1 depicts the block diagram of the proposed pipeline. As illustrated in the figure, it includes four major steps, namely, data pre-processing, feature selection, learning individual SSL models, and finally consolidating these models to produce a combined classifier.

In the first step, the genomic input data are parsed and stored into feature matrices and survival times are retrieved from the clinical records. Since we are dealing with a classification task, the positive class includes those patients who lived above some threshold of interest (e.g., 5 years for the kidney cancer patients) or are still alive and their survival period as of the last follow-up visit passed that threshold. Similarly, the negative class comprises those samples who deceased before the threshold. For future evaluations and training of the pipeline, we randomly select a subset of the labeled samples (15%) and split the rest to five folds to carry out five-fold cross-validation according to a specified seed that we change for repeated runs of the same experiments.

In the next step, to avoid overfitting, it is necessary to perform a feature reduction step. We chose the minimum redundancy maximum relevance (mRMR) [13] feature selection technique which has been successfully applied in a number of studies [13], [20], [21], [18]

that deal with high dimensional data from high-throughput experiments. mRMR is an incremental search algorithm that looks for the subset of features with highest relevance to the classes and lowest redundancy. Both relevance and redundancy are defined in terms of the mutual information. Specifically, relevance is defined as the average mutual information between the selected feature subset and the classes. Redundancy, on the other hand, is defined as the average pairwise mutual information between all the samples in the subset. Clearly in the ideal case where all the selected features are independent, the pairwise mutual information between any two features is zero and as a result, the redundancy measure will be zero. Since, this is a multi-objective optimization and there could be many non-dominated solutions, a single objective function is defined as a weighted sum of relevance and negative redundancy. As our goal here is to select those features that are the key determinants of the long-term vs the short-term survival, we computed the z-score of each modality input matrix and discretized its entries into three levels corresponding to those feature values that fall within the ranges $(-\infty, 1.5\sigma]$, $(-1.5\sigma, 1.5\sigma)$ and $[1.5\sigma, \infty)$. Clearly, the set of differentially expressed genes (DEGs) are expected to fall within the two extreme ranges, whereas those genes which are uncorrelated to the clinical endpoint, are assumed to exhibit smaller variability across the positive and the negative classes. To find the optimal size of selected feature sets as well as the pipeline hyper-parameters (e.g. ambient and intrinsic space regularizers, γ_I and γ_A) we used the validation set and conducted a grid search to pick the best configurations.

Once the number of features have been reduced, we train multiple Laplacian SVM classifiers, one for each modality. We used the library developed by the authors in [2] to train our models. The heat (Gaussian) affinity kernel was selected to generate the edge weights and Euclidean distance function to compute the distance between the nodes. We used 5 nearest neighbor method to filter out the links between distant nodes. Moreover, we selected polynomial kernel of degree three (cf. equation 2) which turns out to strike a good balance between the generalizability and the expressivity of the individual models. Lastly, to come up with a consolidated model, we adopted the stacked generalization strategy [25] to weigh each of the individual models according to their prediction scores. This second layer model is trained over the prediction scores produced by the single modality sub-models using a linear-kernel support vector machine and thereby the prediction scores from the previous layer are normalized so that no single modality model dominates the role of other models just by having a higher range of prediction scores.

IV. RESULTS

We used the proposed pipeline to predict survival of the neuroblastoma (NB) and the kidney (KIRC) cancer patients. In this section, we address two critical questions: (1) Does the SSL approach improve prediction performance? And (2) How does the overall pipeline perform relative to the individual data models (denoted by Model #1, ..., #n in Figure 1)?

A. Supervised vs. semi-supervised strategies

We compared the performance of the Laplacian SVM classifier (LapSVM) to its supervised counterpart, the support vector machines (SVM). We evaluated the performance of each

model 100 times using different cross-validation training, test, and, validation sets. Figure 2 depicts the box plot of the resulting performance measures, for both the neuroblastoma and kidney datasets and the three RNA-seq modalities mentioned above. Interestingly, for the NB dataset, the trained models (both supervised and semi-supervised) can accurately predict the 9-year survival with the semi-supervised models doing significantly better ($p = 0.0, 0.0, 0.0$ for gene, junction, and isoform, respectively). For the KIRC dataset, although the LapSVM outperforms the supervised SVM on average, the performance improvement is not statistically significant for the junction and isoform modalities ($p = 3.5e - 4, 0.2, 0.15$ for gene, junction, and isoform, respectively.) This is because of the more heterogeneous nature of the kidney cancer and the low predictivity of RNA-seq for this dataset, as evidenced by the lower prediction accuracies compared to the neuroblastoma dataset.

To make the role of unlabeled data more tangible, in another series of experiments, we trained the semi-supervised LapSVM models with and without the unlabeled data. In other words, we would like to uncover the impact of adding the unlabeled data on the prediction accuracy. To this end, we repeated the same set of experiments as in the previous subsection but this time with an additional semi-supervised model that does not leverage the unlabeled samples (LapSVM L). Figure 2 shows the resulting accuracies.

According to the figure, adding the unlabeled samples significantly improves the prediction accuracy. In fact when unlabeled samples are not used, the average performance falls below that of the supervised SVM models.

B. Semi-supervised multi-modal pipeline vs. individual data models

Finally, we investigated whether the semi-supervised multi-modal pipeline performs at least as best as the individual sub-models. Moreover, we computed similar results for the supervised counterpart of the proposed pipeline, that is, when the semi-supervised based single modality models are replaced with the supervised SVMs. Table III illustrates the mean along with the standard variation of the prediction accuracies across 100 randomly initialized settings both for the individual models and the pipelines. According to the Table a few key points become evident. First, the semi-supervised models consistently surpass the supervised ones and, second, by using a stacking strategy the prediction accuracy increases synergistically (as is the case for the KIRC dataset) or leads to performances that closely follow the best sub-models (as is the case with the NB dataset.) According to the table the combined model predicts the sample labels with higher or competitive accuracies compared to the individual models.

V. DISCUSSION

Prediction of survival is a critical step in decision making for patient's therapy. A wrong estimate on patient's survival may lead to a choice which undermines the quality of life or the success of the selected treatment. The current study aims to help doctors to make more objective estimates based on the genomic data recorded during the course of patient's treatment. In so doing, we developed a pipeline that predicts cancer survival using multiple modalities of high-dimensional transcriptomic data. We showed that by exploiting the manifold structure of the cancer input sources and integrating multiple modalities into a

semi-supervised learning framework we can achieve high accuracy. We showed that omics data alone can explain a significant portion of the clinical outcome.

While we designed this pipeline for the task of cancer survival prediction, its application is not limited to this area. In fact, such a pipeline is applicable to other domains where there are not enough labeled data available and plenty of unlabeled data exist.

In this study, we only focused on transcriptomic data modalities. However, a more comprehensive approach may involve integration of different levels of genomic data. In fact, cancer has been known to result from dysregulation by multiple molecular mechanisms [7], [11] which can manifest itself in changes in the DNA structure, copy number, DNA methylation, histone modification, and miRNA regulation. As a result, no single level of genomic data by itself can explain the outcome independently. Therefore, it may be beneficial to integrate the results of other high-throughput experiments in different genomic levels to boost the prediction accuracy which highlights the role of big-data analytics. The integration of different layers of genomic data using SSL has been explored in some of the recent published works such as [14], [22]. Another direction that may be worth following is the integration of both clinical and omics data. Clinical prognostic factors convey valuable information that is not available in omics data. Even if this information exists, we may not be able to capture it or effectively process it in the pre-processing and feature reduction phases. Finding an effective way to combine these two inherently different sources of data may be interesting future research.

REFERENCES

- [1]. Belkin M and Niyogi P Semi-supervised learning on riemannian manifolds. *Machine learning*, 56(1–3):209–239, 2004.
- [2]. Belkin M, Niyogi P, and Sindhvani V Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of Machine Learning Research*, 7:2399–2434, 2006.
- [3]. Chakraborty D and Maulik U Identifying cancer biomarkers from microarray data using feature selection and semisupervised learning. 2014.
- [4]. Chapelle O, Scholkopf B, and Zien A Semi-supervised learning. 2006.
- [5]. Chapelle O, Scholkopf B, Zien A, et al. Semi-supervised learning. 2006.
- [6]. Chen X and Yan GY Semi-supervised learning for potential human microrna-disease associations inference. *Scientific Reports*, 4, 2014.
- [7]. Chin L and Gray JW Translating insights from the cancer genome into clinical practice. *Nature*, 452(7187):553–563, 2008. [PubMed: 18385729]
- [8]. Cui Y, Cai XD, and Jin Z Semi-supervised classification using sparse representation for cancer recurrence prediction. 2013 IEEE International Workshop on Genomic Signal Processing and Statistics (Gensips 2013), pages 102–105, 2013.
- [9]. Gripp S, Moeller S, Bolke E, Schmitt G, Matuschek C, Asgari S, Asgharzadeh F, Roth S, Budach W, Franz M, and Willers R Survival prediction in terminally ill cancer patients by clinical estimates, laboratory tests, and self-rated anxiety and depression. *J Clin Oncol*, 25(22):3313–20, 2007. [PubMed: 17664480]
- [10]. Gui J, Wang SL, and Lei YK Multi-step dimensionality reduction and semi-supervised graph-based tumor classification using gene expression data. *Artif Intell Med*, 50(3):181–91, 2010. [PubMed: 20599367]
- [11]. Hanash S Opinion - integrated global profiling of cancer. *Nature Reviews Cancer*, 4(8):638–644, 2004. [PubMed: 15286743]

- [12]. Hassanzadeh HR, Phan JH, and Wang MD A semi-supervised method for predicting cancer survival using incomplete clinical data. In Engineering in Medicine and Biology Society, 2015. EMBS 2015. 37th Annual International Conference of the IEEE, page TBD. IEEE.
- [13]. Huang T, Chen L, Cai Y-D, and Chou K-C Classification and analysis of regulatory pathways using graph property, biochemical and physicochemical property, and functional property. PLoS One, 6(9):e25297, 2011. [PubMed: 21980418]
- [14]. Kim D, Joung JG, Sohn KA, Shin H, Park YR, Ritchie MD, and Kim JH Knowledge boosting: a graph-based integration approach with multi-omics data and genomic knowledge for cancer clinical outcome prediction. J Am Med Inform Assoc, 22(1):109–20, 2015. [PubMed: 25002459]
- [15]. Kim D, Shin H, Sohn KA, Verma A, Ritchie MD, and Kim JH Incorporating inter-relationships between different levels of genomic data into cancer clinical outcome prediction. Methods, 67(3):344–53, 2014. [PubMed: 24561168]
- [16]. Kim D, Shin H, Song YS, and Kim JH Synergistic effect of different levels of genomic data for cancer clinical outcome prediction. J Biomed Inform, 45(6):1191–8, 2012. [PubMed: 22910106]
- [17]. Kim D and Shin H Breast cancer survival prediction using labeled, unlabeled, and pseudo-labeled patient data. Journal of the American Medical Informatics Association, 20(4):613–618, 2013. [PubMed: 23467471]
- [18]. Li B-Q, Huang T, Liu L, Cai Y-D, and Chou K-C Identification of colorectal cancer related genes with mrmr and shortest path in protein-protein interaction network. PloS one, 7(4):e33393, 2012. [PubMed: 22496748]
- [19]. Maulik U, Mukhopadhyay A, and Chakraborty D Gene-expression-based cancer subtypes prediction through feature selection and transductive svm. IEEE Transactions on Biomedical Engineering, 60(4):1111–1117, 2013. [PubMed: 23095982]
- [20]. Mohabatkari H, Beigi MM, Abdolahi K, and Mohsenzadeh S Prediction of allergenic proteins by means of the concept of chou's pseudo amino acid composition and a machine learning approach. Medicinal Chemistry, 9(1):133–137, 2013. [PubMed: 22931491]
- [21]. Mohabatkari H, Mohammad Beigi M, and Esmaeili A Prediction of gabaa receptor proteins using the concept of chou's pseudo-amino acid composition and support vector machine. J Theor Biol, 281(1):18–23, 2011. [PubMed: 21536049]
- [22]. Park C, Ahn J, Kim H, and Park S Integrative gene network construction to analyze cancer recurrence using semi-supervised learning. PLoS One, 9(1):e86309, 2014. [PubMed: 24497942]
- [23]. Scholkopf B and Smola AJ Learning with kernels: Support vector machines, regularization, optimization, and beyond. MIT press, 2002.
- [24]. Shi M and Zhang B Semi-supervised learning improves gene expression-based prediction of cancer recurrence. Bioinformatics, 27(21):3017–3023, 2011. [PubMed: 21893520]
- [25]. Wolpert DH Stacked generalization. Neural networks, 5(2):241–259, 1992.
- [26]. Zhang W, Yu Y, Hertwig F, Thierry-Mieg J, Zhang W, Thierry-Mieg D, Wang J, Furlanello C, Devanarayan V, Cheng J, et al. Comparison of rna-seq and microarray-based models for clinical endpoint prediction. Genome biology, 16(1):1–12, 2015. [PubMed: 25583448]
- [27]. Zhou D, Bousquet O, Lal TN, Weston J, and Scholkopf B Learning with local and global consistency. Advances in neural information processing systems, 16(16):321–328, 2004.

Figure 1:
Pipeline for Semi-Supervised Multi-Modal Prediction Modeling

(a) Neuroblastoma

(b) Kidney renal cancer

Figure. 2:
Comparison of SVM and LapSVM

TABLE I:

Data description

Cancer type	Platform	Modality	#Features
NB		Isoform	263547
		Gene	60781
		Junction	340417
KIRC	Illumina HiSeq 2000	Isoform	73599
		Gene	20531
		Junction	249567

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE II:

Number of labeled vs. unlabeled samples

Cancer type	#Labeled Samples		#Unlabeled Samples	%Unlabeled
	#Positive	#Negative		
NB (9-y survival)	114	105	279	56%
KIRC (5-y survival)	110	140	278	53%

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE III:

Single modality vs integrated model: accuracy comparison

Cancer type	Modality	Mean ACC ($\pm\sigma$): Sup	Mean ACC ($\pm\sigma$): Semi-Sup
NB	Isoform	82.25% (± 2.60)	85.98% (± 1.96)
	Gene	83.04 (± 2.44)%	85.84% (± 2.05)
	Junction	84.50 (± 2.48)%	87.16% (± 1.87)
	Combined	84.93 (± 1.79)%	86.83% (± 1.89)
KIRC	Isoform	63.52 (± 3.10)%	65.02% (± 2.94)
	Gene	61.68 (± 3.07)%	64.67% (± 2.74)
	Junction	62.65 (± 3.19)%	64.68 % (± 2.90)
	Combined	66.07% (± 2.54)	66.20% (± 2.60)