
Encoding Domain Knowledge in Multi-view Latent Variable Models: A Bayesian Approach with Structured Sparsity

Arber Qoku

German Cancer Research Center (DKFZ)
German Cancer Consortium (DKTK)
Goethe University Frankfurt, Germany
arber.qoku@dkfz.de

Florian Buettner

German Cancer Research Center (DKFZ)
German Cancer Consortium (DKTK)
Frankfurt Cancer Institute, Germany
Goethe University Frankfurt, Germany
florian.buettner@dkfz.de

Abstract

Many real-world systems are described not only by data from a single source but via multiple data views. In genomic medicine, for instance, patients can be characterized by data from different molecular layers. Latent variable models with structured sparsity are a commonly used tool for disentangling variation within and across data views. However, their interpretability is cumbersome since it requires a direct inspection and interpretation of each factor from domain experts. Here, we propose MuVI, a novel multi-view latent variable model based on a modified horseshoe prior for modeling structured sparsity. This facilitates the incorporation of limited and noisy domain knowledge, thereby allowing for an analysis of multi-view data in an inherently explainable manner. We demonstrate that our model (i) outperforms state-of-the-art approaches for modeling structured sparsity in terms of the reconstruction error and the precision/recall, (ii) robustly integrates noisy domain expertise in the form of feature sets, (iii) promotes the identifiability of factors and (iv) infers interpretable and biologically meaningful axes of variation in a real-world multi-view dataset of cancer patients.

1 INTRODUCTION

In many real-world applications, complex systems are characterized via multiple data views. That is, observations are represented by multiple groups of distinct features. These




Figure 1: Performance of our model against competitive models. MuVI achieves a higher F_1 score for recovering true active features while maintaining a low reconstruction error, even when using an uninformative prior (MuVI_U).

groups of features often describe different and complementary input sources that are required for a comprehensive characterization of a sample. For example, in genomic medicine, a single patient can be described by quantifying different molecular layers such as the proteome, the microbiome and the transcriptome.

Latent variable models are powerful statistical tools that uncover the axes of variation between samples and data views, by inferring unobserved hidden states from the observable high-dimensional data. To disentangle the sources of heterogeneity driving intra-view and inter-view variation in a meaningful manner, it is key that latent variable models are expressive and interpretable. Modeling approaches with high expressive power are based on autoencoders (Ainsworth et al., 2018) or Gaussian Process Latent Variable models (Damianou et al., 2012). How-

ever, the non-linearities introduced to model complex dependencies hinder interpretability to the extent that linear approaches remain the dominant tool for analyzing multi-view data (Argelaguet et al., 2018). Recently proposed AE-based approaches attempt to balance expressive power and interpretability by combining non-linear encoders with linear decoders (Svensson et al., 2020; Lotfollahi et al., 2023). Commonly used approaches that yield an interpretable decomposition of multi-view data are factor analysis models with structured sparsity on the view, factor and feature level. This modeling task is challenging since the priors inducing the structured sparsity should (i) be amenable to efficient inference strategies that scale to large datasets, (ii) identify true active features and factors across views with a high precision and recall, and (iii) faithfully model the data with a low reconstruction error. Current state-of-the-art methods typically achieve structured sparsity by incorporating a spike-and-slab prior (Argelaguet et al., 2018) or a three-parameter beta (TPB) prior (Zhao et al., 2016). While these models result in sparse solutions, a careful inspection of each latent factor is necessary to provide a meaningful interpretation that is relevant to the domain of study. This manual annotation is cumbersome, time consuming and requires highly-specialized expertise. In addition, the inferred latent representations cannot be directly compared across several training instances as these approaches are unidentifiable unless additional constraints are put on their latent components (Anderson and Rubin, 1956). In brief, non-identifiability in factor analysis states that multiple solutions can produce identical likelihoods, e.g. by applying orthogonal transformations on the latent components, or by permuting the factor indices.

To automate a consistent annotation process, we propose to leverage the partial knowledge on the structure of the latent space that is readily available in many domains. For example, in genomic medicine, disease states are often characterized by sets of features (or pathways), comprising genes that are known to act in a coordinated manner. These data, collected over many years of scientific progress, are curated in dedicated pathway databases. Integrating this domain knowledge in a principled manner as sparsity priors in multi-view models is challenging. Domain knowledge is noisy (e.g. pathways contain many false positive and false negative annotations) and domain knowledge is incomplete and often only available for a subset of views (e.g. in genomic medicine pathway information is only readily available for the transcriptome and the genome).

In this work we propose a novel **multi-view** latent variable model with domain-informed structured sparsity (MuVI) for incorporating domain knowledge via structured sparse priors, and analyzing multi-view data in an inherently explainable manner. We first introduce a modified horseshoe prior for inducing structured sparsity that (i) maintains a low reconstruction error while identifying true active features more reliably than state-of-the-art sparse priors, and

(ii) thereby facilitates the integration of domain knowledge from noisy feature sets. Our domain-informed priors render the inferred latent variables directly interpretable without any further interaction with domain expert, by tagging each latent variable with its corresponding feature set (or pathway in genomic medicine). Briefly, the contributions of our work are as follows:

- We demonstrate in a comprehensive evaluation that our model outperforms state-of-the-art approaches in terms of the reconstruction error and precision, recall and F_1 score.
- We show that our model utilizes prior information efficiently by recovering correct signals from noisy feature sets and is robust against poorly specified priors.
- We demonstrate the practical utility of our model on a large multi-view dataset of cancer patients by inferring interpretable and biologically meaningful axes of variation.

2 RELATED WORK

Factor analysis (FA) is a fundamental approach for estimating and understanding the correlation structure among observed variables (Thurstone, 1931), which has inspired the development of numerous latent variable models. Due to its simplicity, however, standard FA is unable to model observations from multiple sources. Extensions, such as the canonical correlation analysis (CCA) (Hotelling, 1992; Klami et al., 2013), or the group factor analysis (GFA) (Klami et al., 2014) model paired observations simultaneously by learning linear dependencies underlying two or more data sources. A central assumption of the GFA is that the multi-view observations are a manifestation of a lower dimensional common latent space corrupted by Gaussian noise. The main components that describe this relationship are the factor-to-feature linear mappings or the so-called factor loadings. Factor loadings encode the structure of each factor, and play an important role in the interpretation of the model. Hence, factor loadings are suitable for introducing statistical assumptions on their underlying structure such as sparsity. For instance, GFA quantifies the association between view and factor by extending the automatic relevance determination (MacKay et al., 1994).

2.1 Sparsity Inducing Priors

In non-Bayesian approaches, sparsity is commonly handled by introducing additional terms to the optimization objective posed on the model parameters. A canonical example is the L1 penalty or lasso (Tibshirani, 1996). On the other hand, Bayesian approaches achieve sparse solutions via sparsity inducing priors. For instance, the double exponential, or Laplace prior is the Bayesian counterpart of the

lasso (Park and Casella, 2008). Another popular choice is the discrete spike-and-slab prior (Mitchell and Beauchamp, 1988), a mixture of a Dirac delta distribution centered at zero for pruning irrelevant signals, and a normal distribution for modeling larger signals. Recently, the spike-and-slab lasso (SSL) (Ročková and George, 2018) has emerged as a combination of two Laplace distributions that are parameterized to emulate both the spike and the slab component. Other flavors of shrinkage priors (Polson and Scott, 2010) include the horseshoe prior (Carvalho et al., 2009), which offers a continuous relaxation of the spike-and-slab approach, thereby preserving sparse properties while providing computational benefits such as being differentiable.

2.2 Structured Sparsity in Multi-view LVMs

Several Bayesian approaches have successfully utilized the synergy between latent variable models and sparsity inducing priors (Bernardo et al., 2003; Engelhardt and Stephens, 2010; Lan et al., 2014). In the multi-view setting, (Zhao et al., 2016) propose a hierarchical Bayesian GFA with structured sparsity, facilitated by a cascading three parameter beta prior (Armagan et al., 2011). The consequences of this addition are twofold, the column-wise sparsity supports the inference of associations between views, while element-wise sparsity encourages feature selection within individual factors. Multi-omics factor analysis (MOFA) (Argelaguet et al., 2018) assumes a similar structured sparsity in two levels, which is achieved by combining an automatic relevance determination (ARD) prior (MacKay et al., 1994) with a spike-and-slab prior. Beyond linear models, there is an array of non-linear modeling approaches that generalize the ideas of linear LVMs and sparse priors to more complex data settings. These include approaches based on variational autoencoder (Ainsworth et al., 2018) or Gaussian Process Latent Variable Models (Damianou et al., 2012). Such approaches, however, are difficult to interpret in practice and have therefore not received widespread application. For example, in the latter GP-based approach, loading matrices are marginalized out, and cannot be used to identify any physical or biological processes captured by individual factors. Recent hybrid approaches (Svensson et al., 2020; Lotfollahi et al., 2023) combine non-linear encoders with linear decoders to maintain the expressive power of autoencoders while preserving the interpretability of linear models. While these approaches were proposed for single data views, extension to the multi-view setting is straightforward.

2.3 Integration of Domain Knowledge in LVMs

Typically, before training sparse latent variable models, the prior is set to be relatively uninformative, allowing only the observations to mold the posterior. An alternative approach is introducing an informed structured sparsity into the model based on external expertise in the domain of in-

terest. The factorial single-cell latent variable model (Buettner et al., 2017) attempts to bridge this gap by explicitly integrating, and jointly modeling a collection of domain relevant feature sets such as gene set annotations. The model extends the spike-and-slab to be informed by the presence or absence of a feature, and infers pre-labeled latent axes of variation. However, this method can only handle a single data view and the proposed inference scheme does not scale to large sample sizes. An orthogonal approach for identifying latent variables corresponding to pathways is the single sample gene set enrichment analysis (ssGSEA) (Subramanian et al., 2005; Barbie et al., 2009), which computes a sample-wise enrichment of a pathway. However, this approach does not account for noisy annotations, treats pathways as independent and can only be applied to a single data view. A multi-view approach with structured sparsity that admits prior information in terms of feature sets, and handles partial and potentially noisy priors has not yet been explored.

3 DOMAIN-INFORMED MULTI-VIEW MODELING

3.1 Background and Notation

Let $\mathbf{y}_i \in \mathbb{R}^D$ denote a D -dimensional observation for $i \in \{1, \dots, N\}$, and $G_m \subseteq \{1, \dots, D\}$ describe a mutually disjoint grouping of the features into M data views, where $G_p \cap G_q = \emptyset$ for $p \neq q \in \{1, \dots, M\}$. For simplicity, we assume a permutation of the features D such that the first D_1 features belong to G_1 , the second D_2 features to G_2 and so on. As a result, we may rewrite the collection of the observations as a matrix $\mathbf{Y} \in \mathbb{R}^{N \times D}$ comprising M coupled views $[\mathbf{Y}^{(1)}, \mathbf{Y}^{(2)}, \dots, \mathbf{Y}^{(M)}]$, where $\mathbf{Y}^{(m)} \in \mathbb{R}^{N \times D_m}$. The main goal is to then represent each observation \mathbf{y}_i in terms of a low-dimensional set of latent factors $\mathbf{x}_i \in \mathbb{R}^K$, $K \ll D$. The relationship between the observations and the latent factors is described by a set of view-specific factor loadings $\mathbf{W}^{(m)} \in \mathbb{R}^{D_m \times K}$. Then, the recipe for the group factor analysis framework posits the following generative process to the observed data,

$$\mathbf{y}_i^{(m)} \sim \mathcal{N}(\mathbf{W}^{(m)} \mathbf{x}_i, \mathbf{\Psi}^{(m)}) \quad (1)$$

where each latent variable typically follows an isotropic standard normal distribution

$$\mathbf{x}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (2)$$

The residuals are denoted by $\mathbf{\Psi}^{(m)} = \text{diag}(\sigma_j^{2(m)})$, a diagonal matrix storing the marginal variances $\sigma_j^{2(m)}$ of each variable j in view m . Due to conjugacy properties, setting

$$\sigma_j^{2(m)} \sim \Gamma^{-1}(\alpha_\sigma, \beta_\sigma) \quad (3)$$

is a common choice, where $\Gamma^{-1}(\alpha, \beta)$ describes the inverse-Gamma with shape and scale parameters α and β .

Finally, an important component of GFA is the collection of the factor loadings $\mathbf{W}^{(m)}$ which applies a linear projection of the latent variable \mathbf{x}_i to $\mathbf{y}_i^{(m)}$. In order to facilitate the interpretability of the factor-to-feature mapping, (Klami et al., 2014) propose a structured sparsity for the columns of $\mathbf{W}^{(m)}$, such that each factor falls into two distinct categories. A non-zero factor loading vector $\mathbf{w}_k^{(m)}$ indicates an active factor k , while $\mathbf{w}_k^{(m)} = \mathbf{0}$ indicates an inactive factor k in view m . Hence, a factor is either shared across an arbitrary subset of views, or private to a specific view.

3.2 MuVI

We follow a similar approach, inspired by the success of the horseshoe prior (Carvalho et al., 2009, 2010), and introduce a *view-factor-local* shrinkage prior on the factor loadings to enable both column-wise and element-wise shrinkage:

$$w_{j,k}^{(m)} \sim \mathcal{N}\left(0, \left(\tau^{(m)} \delta_k^{(m)} \lambda_{j,k}^{(m)}\right)^2\right), \quad (4)$$

where each scale in the hierarchy follows a positive Cauchy distribution,

$$\tau^{(m)} \sim \mathcal{C}^+(0, 1) \quad (5)$$

$$\delta_k^{(m)} \sim \mathcal{C}^+(0, 1) \quad (6)$$

$$\lambda_{j,k}^{(m)} \sim \mathcal{C}^+(0, 1). \quad (7)$$

Each level in the hierarchy contributes to the overall structured sparsity of the factor loadings. In particular, $\delta_k^{(m)}$ serves as an automatic relevance determination (ARD) mechanism (MacKay et al., 1994) for factor k in view m , effectively decoupling this factor from the rest of the views. At the same time, $\lambda_{j,k}^{(m)}$ acts as a regulator on each individual loading, encouraging the model to seek simpler solutions that describe each factor in terms of fewer features. However, due to the heavy tails of the Cauchy distribution, weakly identified loadings under the horseshoe prior can easily escape the regularization. To counteract this behavior, the regularized horseshoe guarantees a non-zero penalty even for large weights (Piironen and Vehtari, 2017). Concurrently, the regularized parameterization helps integrate prior information about the structure of the latent factors into the model. We update Equation 4 as follows. To simplify notation, we sometimes drop the view-specific superscript m , which can also be implicitly encoded in j . Let $\gamma_{j,k} = \tau \delta_k \lambda_{j,k}$, then

$$w_{j,k} \sim \mathcal{N}\left(0, \frac{(c_{j,k} \gamma_{j,k})^2}{c_{j,k}^2 + \gamma_{j,k}^2}\right) \quad (8)$$

Note that in Equation 8 the $c_{j,k}^2$ parameter relates to a specific factor loading rather than being a global parameter as originally defined in the regularized horseshoe. For a relatively large $c_{j,k}^2 \gg \gamma_{j,k}^2$, $w_{j,k}$ is scaled by a factor of

nearly $\gamma_{j,k}$, rendering the effect of the additional penalty insignificant, and reinstating the original horseshoe prior. On the other hand, when $c_{j,k}^2 \ll \gamma_{j,k}^2$, $w_{j,k}$ is scaled by a factor of nearly $c_{j,k}$, thus setting an upper bound on corresponding factor loadings. Equivalently, the regularized horseshoe can be seen as a continuous alternative to the discrete spike-and-slab prior (Mitchell and Beauchamp, 1988) with a finite slab width. Piironen and Vehtari (2017) suggest assigning an inverse-Gamma distribution as a weakly informative prior to

$$c_{j,k}^2 \sim \Gamma^{-1}(\alpha_c, \beta_c). \quad (9)$$

For all our experiments we opt for $\alpha_c = \beta_c = 0.5$, as this encourages sparsity while still allowing strong signals to escape the regularization due to the heavy right tail.

3.3 Integrating Prior Knowledge from Noisy Feature Sets

Next, we attempt to integrate prior knowledge in terms of noisy feature sets into our model.

Feature sets A feature set \mathbf{I}^D is a collection of binary variables $I_j \in \{0, 1\}$, $j \in \{1, \dots, D\}$, where $I_j = 1$ indicates the presence of feature j , and $I_j = 0$ its absence. Assume we have substantial knowledge about the underlying structure of the factor loadings in terms of feature sets. That is, for every latent dimension k and a set of features D we are given a corresponding feature set \mathbf{I}_k^D , such that $|w_{j,k}| > 0$ if $I_{j,k} = 1$, and $w_{j,k} = 0$ otherwise. Each latent factor can be seen as a factor-to-feature mapping of the present features determined by the feature set. In terms of Equation 8, this can also be achieved by setting the corresponding slab width $c_{j,k} \approx 0$, i.e. applying an infinitely large weight decay penalty to $w_{j,k}$.

Noisy feature sets In practice, we rarely have access to such pristine ground truth. However, in some cases, we may exploit existing domain knowledge to develop a prior belief about a plausible structure of the factor loadings by accommodating noisy feature sets into our modeling approach. Let $\tilde{\mathbf{I}}^D$ be a noisy version of a feature set \mathbf{I}^D , where a subset of the binary variables $I_Q, Q \subset \{1, \dots, D\}$ has been flipped to generate $\tilde{\mathbf{I}}^D$, inserting a non-zero fraction of false positives and false negatives. To integrate noisy feature sets into MuVI, we relax the hard regularization penalty induced by a pre-defined $c_{j,k}$, allowing the adaptation of incorrect signals, given sufficient evidence from the data. In our approach, we opt for an auxiliary hyperparameter $0 < \alpha_{j,k} \leq 1$ which scales $c_{j,k}$. A value of $\alpha_{j,k} = 1.0$ poses no prior penalty to $w_{j,k}$, while $\alpha_{j,k} < 1.0$ leads to a smaller slab width a priori. Consequently, a smaller $\alpha_{j,k}$ translates to a stronger prior belief. Experiments show that values around $0.01 \leq \alpha_{j,k} \leq 0.05$ for absent features in a prior collection of feature sets consistently provide the best results across different datasets and training scenarios.

Pure variables Present features that depend on only one latent factor are also referred to as pure variables (Bing et al., 2020). Work on the identifiability of factor analysis (Bing et al., 2020; Anderson and Rubin, 1956) proves that the existence of at least two pure variables is a sufficient condition for yielding identifiable solutions. That is, for every factor $k \in \{1, \dots, K\}$, there exist at least two features $j, j' \in \{1, \dots, D\}$ such that $I_{j,k} = 1, I_{j',k} = 1$ for all $j \neq j'$. Under the assumption of known pure variables, MuVI results in provably identifiable factors. However, since we are only given a noisy feature set, we cannot know whether pure variables really are present for all factors. We therefore investigate empirically to what extent our model preserves the identifiability properties when increasing the amount of noise in the prior annotations.

3.4 Inference

The joint model can be written as

$$\begin{aligned} p(\mathbf{Y}, \Theta) &= p(\mathbf{Y}, \mathbf{X}, \mathbf{W}, \Lambda, \Delta, \tau, \mathbf{C}, \Psi) \\ &= p(\mathbf{Y} | \mathbf{X}, \mathbf{W}, \Psi) p(\Psi) \\ &\quad \times p(\mathbf{X}) p(\mathbf{W} | \Lambda, \Delta, \tau, \mathbf{C}) \\ &\quad \times p(\Lambda) p(\Delta) p(\tau) p(\mathbf{C}), \end{aligned} \quad (10)$$

where $\Theta = \mathbf{X}, \mathbf{W}, \Lambda, \Delta, \tau, \mathbf{C}, \Psi$, and

$$\begin{aligned} \mathbf{Y} &= \left\{ y_{i,j}^{(m)} \right\}, \Psi = \left\{ \sigma_j^{(m)} \right\}, \mathbf{X} = \left\{ x_{i,k} \right\}, \mathbf{W} = \left\{ w_{j,k}^{(m)} \right\}, \\ \Lambda &= \left\{ \lambda_{j,k}^{(m)} \right\}, \Delta = \left\{ \delta_k^{(m)} \right\}, \tau = \left\{ \tau^{(m)} \right\}, \mathbf{C} = \left\{ c_{j,k}^{(m)} \right\}. \end{aligned}$$

Inference is performed by introducing a fully factorized family of parameterized distributions $q_\phi(\Theta) = \prod_{\theta \in \Theta} q_\phi(\theta)$ to approximate the intractable posterior $p(\Theta | \mathbf{Y})$. We maximize the evidence lower bound with respect to the variational parameters ϕ , which, in return reduces the gap between the true and the approximate posterior in terms of the KL divergence. The family of normal distributions is a natural choice for approximating \mathbf{X} and \mathbf{W} , whereas, for the rest of the parameters we assert the log-Normal distribution to ensure positive samples (Ghosh et al., 2018). The resulting optimization objective

$$\mathcal{L}(\phi) = \mathbb{E}_{q_\phi} [\log p(\mathbf{Y}, \Theta) - \log q_\phi(\Theta)] \quad (11)$$

is amenable to automated stochastic variational inference (Ranganath et al., 2014; Hoffman et al., 2013), which involves sampling from the variational distribution, and taking unbiased but noisy Monte Carlo estimates of the gradient $\nabla_\phi \mathcal{L}$. In addition, the reparameterization trick (Kingma and Welling, 2013) further stabilizes the optimization procedure, in which the random variables are expressed as a combination of deterministic variables and external random noise, thereby greatly reducing the variance of the MC estimates. We provide an implementation of MuVI using Pyro (Bingham et al., 2019) on GitHub¹.

¹<https://github.com/MLO-lab/MuVI>

4 EXPERIMENTS

We first evaluate different aspects of MuVI empirically on a wide range of simulated settings. Ideally, the model performs well in the following general tasks:

- Learning a meaningful latent representation with low reconstruction loss - both in the presence and absence of prior information.
- Utilizing prior information efficiently to recover correct signals from noisy feature sets, and being robust against poorly specified or entirely incorrect priors.
- Promoting an implicit flow of the prior information from the informed to the uninformed views via the shared factors.

We compare our model against four baselines that admit observations across multiple views, and learn a sparse representation of the common latent space: the Bayesian group factor analysis with structured sparsity (BASS) (Zhao et al., 2016), multi-omics factor analysis (MOFA) (Argelaguet et al., 2018) and a multi-view autoencoder (AE), as a naive multi-view extension of the interpretable autoencoder (Svensson et al., 2020). We also include the group factor analysis (GFA) (Klami et al., 2014) for the synthetic evaluation, while keeping in mind that GFA does not infer sparse solutions regarding the factor loadings. Finally, we demonstrate the practical utility of MuVI on a real-world dataset from The Cancer Genome Atlas (TCGA) (Tomczak et al., 2015).

4.1 Synthetic Experiments

4.1.1 Data Generation

We adopt the data generation process from (Zhao et al., 2016), and compile a synthetic dataset of $N = 200$ samples across four views, each comprising $D_1 = D_2 = D_3 = D_4 = 400$ features. The latent space consists of $K = 15$ factors, that are linearly transformed by a set of sparse factor loadings. Each weight is sampled independently from a standard normal distribution, where loadings with an absolute value of less than 0.1 are set to zero, to emphasize the gap between active and inactive signals. In addition, we randomly set 85%-95% of the loadings to zero. The relationship between the data views is explained by the structured sparsity of the loadings. In particular, we generate all possible combinations of four binary variables, resulting in 15 distinct relationship configurations for each factor as depicted in the top heatmap of Figure 2. A dark entry means the factor does not contribute to the corresponding view, and decouples it from the rest of the views. For instance, the factor with index 0 is fully shared across all views, whereas the factors 7, 11, 13 and 14 explain variability pertaining to each view individually. The rest of the factors




Figure 2: Depiction of cross-view relationships in the synthetic data. True relationship (top) versus inferred relationship based on the factor scales learned by MuVI (bottom).

exhibit a mixture of partially shared configurations. We extract the true feature sets, indicating active and inactive features for each latent dimension, and generate potentially noisy feature sets which serve as prior information during training. We perturb the feature sets by swapping a fraction of the true positives with true negatives. The severity of the noise translates to a poorly specified prior belief, where 100% noise means that the prior belief is entirely incorrect, up to the number of the expected active features. Also, we introduce additional false positive features for factors that are completely inactive, as to further disguise any true underlying relationship within and between views.

4.1.2 Evaluation

We evaluate all models on five randomly generated datasets as described above, and report average metric scores across all views. We assess the quality of the representation in terms of the RMSE between the true observations and the reconstructed views from the common latent space. We quantify the concordance between the inferred factor loadings and the true underlying structure by reporting the precision, recall and the F_1 score. As a prerequisite to calculating the binary scores we extract a binary representation of active and inactive features based on a threshold. For our model and MOFA, the threshold of 0.1 matches the true cutoff between active and inactive signals when generating the sparse loadings. For BASS and AE, the optimal thresholds were 0.05 and 0.01, respectively. In addition, we show the precision and recall curves summarizing results across all possible thresholds in the Appendix A. Finally, we perform factor matching or label switching in a post hoc manner, following (Zhao et al., 2016), where we reorder the factors to match the true order for a direct comparison with the ground truth.

Table 1: Performance comparison on the synthetic data. Average metric scores across five independent runs \pm standard deviation.

| | RMSE | Precision | Recall | F_1 |
|------------------------------------|------------------------------------|------------------------------------|------------------------------------|------------------------------------|
| GFA | 0.324 ± 0.00 | 0.256 ± 0.05 | 0.916 ± 0.05 | 0.402 ± 0.06 |
| BASS | 0.543 ± 0.02 | 0.944 ± 0.05 | 0.899 ± 0.03 | 0.920 ± 0.04 |
| MOFA | 0.331 ± 0.00 | 0.541 ± 0.14 | 0.914 ± 0.02 | 0.672 ± 0.11 |
| AE | 0.325 ± 0.01 | 0.692 ± 0.06 | 0.964 ± 0.01 | 0.805 ± 0.04 |
| MuVI _U (uninformed) | 0.323 ± 0.00 | 0.933 ± 0.05 | 0.995 ± 0.00 | 0.962 ± 0.03 |
| MuVI _{0.5} (1 inf. view) | 0.322 ± 0.00 | 0.966 ± 0.01 | 0.996 ± 0.00 | 0.981 ± 0.00 |
| MuVI _{0.1} (1 inf. view) | 0.322 ± 0.00 | 0.965 ± 0.01 | 0.997 ± 0.00 | 0.982 ± 0.00 |
| MuVI _{0.1} (3 inf. views) | 0.322 ± 0.00 | 0.974 ± 0.01 | 0.998 ± 0.00 | 0.986 ± 0.01 |

4.1.3 Training

The parameter and training settings for the baselines are summarized in the Appendix A. We perform a simple search on the single hyperparameter of MuVI, the auxiliary constant $\alpha_{j,k}$ for the factor loadings that are not part of the prior feature sets. Among $\{0.01, 0.03, 0.05, 0.1\}$, $\alpha_{j,k} = 0.03$ reliably performs the best. A complete sensitivity analysis provided in the Appendix B shows that our approach is robust against the choice of α . We train several models with different noise configuration and prior information availability. We start with an uninformed model and iteratively increase the amount of the prior information by informing a single view, two views and eventually three views. An important consideration when injecting prior information for a subset of views is assuming that the features of the rest of the views are not present in the prior feature sets, and informing them accordingly. In other words, we choose the same $\alpha_{j,k}$ for all the factor loadings for the uninformed views that we choose for the factor loadings of the informed views that are absent in the prior feature sets. The main rationale behind this decision is encouraging the informed views to learn first, and gradually inform the rest of the uninformed views implicitly via the structure of the shared factors. Also, we observed empirically that the uninformed views are unrestricted in terms of learning and typically converge very early during training and are unable to “unlearn” a suboptimal structure of the factor loadings.

4.1.4 Results

Overall performance against baselines We provide a visual summary of the experiments in Figure 1, and a detailed overview of the results in Table 1. In terms of the quality of the latent representation, measured by the RMSE, most models perform comparably well. It is worth noting that the hybrid AE, combining a non-linear encoder with a linear decoder, does not yield a better latent representation. Recent theoretical work on posterior collapse in linear VAEs (Lucas et al., 2019) investigates and proves that there is no benefit in using a non-linear encoder when

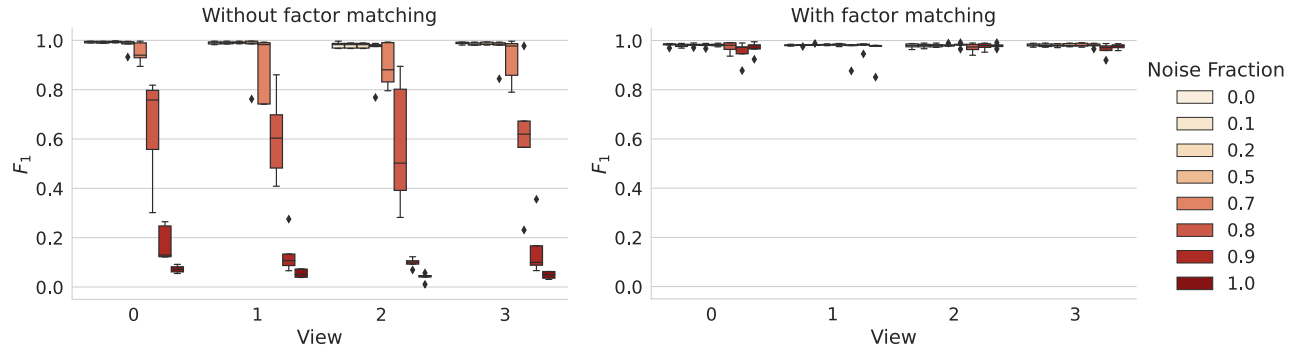


Figure 3: Performance across multiple prior noise configurations when informing only a single view (view 0). Informing a factor in one view is sufficient to promote identifiability in uninformed views even for medium-to-high noise (left). Factor matching is required only for noise levels greater than 70% (right). Each boxplot comprises five independent runs of the same noise fraction.

the decoder is linear. On the other hand, BASS appears to sacrifice reconstruction quality in favor of additional sparsity, which is supported by a high precision and relatively lower recall. As expected, group factor analysis (GFA) is unable to provide sparse solutions. Our modeling approach strikes a better balance between learning a good latent representation with low reconstruction error and pruning away the majority of superfluous connections between the latent factors and the observed features. In addition, the high recall indicates a better capability of MuVI_U to retrieve virtually all positive signals, which results in the highest F_1 score, even in the absence of an informative prior. In addition to the uninformed version of our model (MuVI_U), we report the results for the informed models. For the first model ($\text{MuVI}_{0.5}$), we simulate a realistic scenario where the available information is very limited and noisy. Specifically, we inform only the first view with highly perturbed feature sets with a noise fraction of 50% ($\text{MuVI}_{0.5}$). Compared to the uninformed model, we notice an increase in the precision, while the recall remains intact, thereby further improving the F_1 score. In the next two models, we first reduce the amount of noise to 10% ($\text{MuVI}_{0.1}$), then increase the number of informed views from one to three. We observe additional improvements in all three binary scores, which further emphasizes the benefit of increasing the quality and the quantity of the prior information.

Learning from noisy priors Next, we assess the robustness of MuVI against severe noise, and its capability of utilizing and transferring useful prior information across all views. Figure 3 depicts a comprehensive overview of the results when informing only a single view, while increasing the amount of prior noise. We first assess the extent to which MuVI is able to infer identifiable factors for increasing levels of noise. Fig. 3 (left) shows that up to a noise fraction of around 70%, MuVI exhibits no difficulties in

learning a good representation, and identifies the true underlying structure of the latent factors with a median F_1 score of almost 1.0, not only for the informed view, but also for the uninformed views. This indicates that MuVI is able to promote identifiability across views. Recovering the true structure of the latent factors first becomes challenging when increasing the noise fraction beyond 80%, with a median F_1 score of 0.62 for 80% noise. As expected, a severe perturbation of at least 90% significantly reduces the ability of our model to find the proper ordering of the factors. However, the plot on the right shows that the model merely loses its ability to identify the correct order of the latent dimensions, as the corrected permutation of the inferred factors corresponds to the true structure across all factors and data views. In the Appendix B we provide additional results by considering the noise fraction as a proxy for the number of pure variables available in the prior information. Furthermore, the quality of the representation remains virtually intact. For comparison, the mean RMSE for the trained MuVI models with a noise fraction of 80% or less is 0.322 for each noise level, and for a noise fraction of 90% and 100% we observe only a marginal increase of the RMSE to 0.326 and 0.328, respectively (see Appendix B). These results still outperform the baselines, and further assert that any amount of prior information benefits our model highly, while a (fully) incorrect prior information does not harm its performance.

Learning cross-view relationships Finally, we look at the ability of MuVI to communicate the underlying relationship between the views. The bottom heatmap in Figure 2 corresponds to the MAP estimates of the factor scales $\delta_k^{(m)}$ learned during inference. The model successfully learns to turn off irrelevant factors and highlights cross-view relationships by assigning large positive values to active and shared factors.

4.2 The Cancer Genome Atlas (TCGA)

We investigate a large dataset of comprehensive multi-omic profiling of over 11 thousand samples from 33 cancer types (Tomczak et al., 2015). Each sample comprises sets of features of various sizes across four biologically distinct views: DNA methylation (6,000), mRNA expression (6,000), microRNA expression (728) and reverse phase protein array or RPPA (312). We include a preprocessing step of centering the data and standardizing each view globally due to the large differences in value ranges across views. In addition, samples exhibit partial or complete missingness in individual views, which we accommodate in our probabilistic modeling approach. Prior to training, we derive gene set annotations from curated public databases such as MSigDB (Liberzon et al., 2015), Reactome (Fabregat et al., 2018) and KEGG (Kanehisa and Goto, 2000). A gene set consists of a group of genes that are biologically meaningful when co-expressed under certain conditions, or describe functionally distinct pathways in biological systems. Since the observed features are genes, each gene set acts as a prior feature set for the factor loadings, thereby encoding a well-defined biological pathway into the corresponding latent factor. As informative prior we consider gene set annotations where at least 15 genes were present in the data, resulting in 360 annotations with a median size of 30. Since gene sets encompass only genomic and transcriptomic features, we are limited to informing only the first two views, namely the DNA and the mRNA. We take advantage of the stochastic variational inference algorithm and propagate batches of 1,000 samples during training. The algorithm terminates when the optimization objective no longer improves significantly in consecutive iterations, after a certain number of patience steps. The training converges in less than 12 minutes on a single NVIDIA Quadro RTX 5000 GPU with 16 GB memory.

Global structure First, we assess the quality of the latent space inferred by our model. We systematically evaluate and compare the performance of MuVI and the other baselines in terms of explaining the underlying structure of the data introduced by the different cancer types. For each model, we first infer and extract the latent representation of each observation. Next, we apply a t-distributed stochastic neighbor embedding (t-SNE) approach to further compress the latent space into two dimensions for better visualization (Van der Maaten and Hinton, 2008). In the left-most plot of Figure 4 we display the new embeddings. Our model is able to learn a meaningful structure by grouping samples of a similar type closer together. Subsequently, we apply a K-Means algorithm by fixing the number of clusters to the number of known cancer types (33). Finally, we compare the learned clusters by the K-Means with the true underlying cancer types and compute the cluster purity and the adjusted Rand index (ARI) to measure the similarity between both clusterings. We perform the same procedure

several times and report the average scores in the first two columns of Table 2. The scores reiterate the results obtained in the synthetic evaluation. Overall, MOFA, AE and MuVI perform comparably well by capturing known cancer types. On the other hand, BASS infers a less informative latent space as indicated by the multiple cluster overlaps (see Appendix C), and the lower clustering scores. Importantly, MuVI is inherently interpretable with each learned factor being directly labeled by a biological pathway: our model intrinsically accommodates prior information in terms of feature sets, thereby guiding the inference of pre-tagged axes of variation from domain expertise. In contrast, rendering the baseline methods interpretable requires a cumbersome manual inspection of the loadings of each factor, with subsequent enrichment analysis searching for pathways in which large loadings are over-represented.

Relevant factors Next, we apply a one-vs-rest Wilcoxon rank test to identify the main pathways that drive the heterogeneity of patients with respect to their cancer types. A complete overview of the test results is provided in the Appendix C. In particular, the results hint at melanogenesis (KEGG) and the androgen response (Hallmark) as relevant pathways for melanoma patients and prostate adenocarcinoma patients, respectively. Melanogenesis and the androgen response describe general biological processes. Melanogenesis describes the process of producing the melanin pigments from melanocytes, commonly found in the epidermis and hair follicles. Androgen response orchestrates the activation of male hormone receptors, and plays an important role in the development and progression of prostate cancer (Fujita and Nonomura, 2019). In the middle of Figure 4 we illustrate a scatter plot by mapping our samples onto the latent space spanned by the two axes informed and described by androgen response and melanogenesis. We highlight three cancer types with the same coloring as the previous t-SNE plot, while fading the remaining 30 cancer types to gray. As anticipated, melanogenesis helps identify the only two groups of samples from tissues with skin cutaneous melanoma (SKCM) and uveal melanoma (UVM). Likewise, androgen response assigns significantly higher scores to patients with prostate adenocarcinoma (PRAD). We assess how predictive each latent factor is to its corresponding cancer type by reporting the average area under the ROC curve (AUC) computed for the latent factors assigned to each cancer type in a one-vs-rest fashion. We compare the obtained results with the baselines, by performing a gene set enrichment analysis (Subramanian et al., 2005) on the factor loadings inferred by BASS, MOFA and AE, and match each latent factor to a gene set. Then, we focus on the latent factors that model androgen response and melanogenesis, and measure the one-vs-rest AUC of the factor scores for PRAD and SKCM + UVM, respectively. Since the same pathway can be assigned to multiple latent factors for MOFA and AE, we se-




Figure 4: Results of the TCGA dataset. A t-SNE embedding space learned from the latent representation of MuVI (left). Inferred latent factors that are informed by androgen response and melanogenesis a priori (center). Top 15 features based on the absolute value of their corresponding factor loadings for the shared androgen response factor among mRNA and RPPA (right).

Table 2: Performance comparison on the TCGA dataset. Summary of the cluster purity and the adjusted Rand index (ARI) for measuring the similarity between the known cancer subtypes and the learned clusters by the K-Means algorithm (first two columns). In addition, we report the average AUC scores of individual factors modeling androgen response and melanogenesis (last two columns). For the baselines it is necessary to perform an additional gene set enrichment analysis (GSEA) to label each factor.

| | ARI | Purity | Androgen Resp. (AUC) | Melanogenesis (AUC) |
|------|---------------------|---------------------|----------------------|---------------------|
| BASS | 0.609 ± 0.01 | 0.763 ± 0.01 | 0.553 ± 0.04 | 0.531 ± 0.01 |
| MOFA | 0.658 ± 0.01 | 0.822 ± 0.01 | 0.569 ± 0.04 | 0.528 ± 0.01 |
| AE | 0.662 ± 0.01 | 0.828 ± 0.01 | 0.821 ± 0.05 | 0.724 ± 0.02 |
| MuVI | 0.660 ± 0.01 | 0.825 ± 0.01 | 0.910 ± 0.03 | 0.872 ± 0.02 |

lect the best performing factor and report the AUC scores in the last two columns of Table 2. MuVI achieves the highest AUC, while rendering each latent factor directly interpretable due to the encoded pathway knowledge a priori.

Factor inspection Finally, we focus on the composition of the inferred latent factors. In the rightmost plots, we inspect the top 15 features of the androgen response factor, which is active and shared among the mRNA and the RPPA view. We decide whether a factor is active in view m if its corresponding factor scale is not virtually zero, i.e. $\delta_k^{(m)} > 0.01$, as well as if the variance explained by the factor exceeds a threshold of $R^2 > 0.5\%$. We reiterate that we informed only mRNA among the two views. Hence, we expect to see annotated genes weighted highly based on the absolute value of their corresponding fac-

tor loadings. A powerful feature of MuVI is the refinement of provided feature sets by adding or removing features in a data-driven manner. Investigating inferred genes for the androgen response factor, we discover several previously known biomarkers of clinical relevance such as STEAP1 (Khanna et al., 2021), SLC45A3 (Perner et al., 2013) and ACPP (Kong and Byun, 2013), which are over-expressed in patients with prostate cancer. We now move to the uninformed view, where we expect the androgen response factor of mRNA to influence the structure of the factor loadings in the RPPA, such that the inferred factor reveals protein markers linked to prostate cancer. Indeed, the inferred BCL2 is an oncoprotein that inhibits the process of apoptosis (Lin et al., 2007), whereas targeting the checkpoint kinase 1 (CHEK1) leads to favorable clinical outcomes for prostate cancer patients (Drápela et al., 2020).

5 CONCLUSION

In this contribution, we have addressed the task of modeling multi-view data with interpretable latent variable models. We use a Bayesian approach with structured sparsity to encode domain knowledge in the priors of a factor analysis model. Our model is able to integrate noisy domain expertise and results in factors that are inherently interpretable via their pre-defined feature sets. In contrast to other sparse multi-view models, our approach is able to recover correct signals from noisy feature sets, while maintaining a low reconstruction error. We demonstrate in a real-world application that our model is able to infer biologically meaningful and clinically relevant subpopulations of cancer patients. Moreover, the ability of MuVI to transfer prior knowledge from informed to uninformed views via the shared factors shows the potential of discovering novel feature sets in related domains.

Bibliography

- Samuel K Ainsworth, Nicholas J Foti, Adrian KC Lee, and Emily B Fox. oi-vae: Output interpretable vaes for non-linear group factor analysis. In *International Conference on Machine Learning*, pages 119–128. PMLR, 2018.
- TW Anderson and Herman Rubin. Statistical inference in. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability: Held at the Statistical Laboratory, University of California, December, 1954, July and August, 1955*, volume 1, page 111. Univ of California Press, 1956.
- Ricard Argelaguet, Britta Velten, Damien Arnol, Sascha Dietrich, Thorsten Zenz, John C Marioni, Florian Buettner, Wolfgang Huber, and Oliver Stegle. Multi-omics factor analysis—a framework for unsupervised integration of multi-omics data sets. *Molecular systems biology*, 14(6):e8124, 2018.
- Artin Armagan, Merlise Clyde, and David Dunson. Generalized beta mixtures of gaussians. *Advances in neural information processing systems*, 24, 2011.
- David A Barbie, Pablo Tamayo, Jesse S Boehm, So Young Kim, Susan E Moody, Ian F Dunn, Anna C Schinzel, Peter Sandy, Etienne Meylan, Claudia Scholl, et al. Systematic rna interference reveals that oncogenic kras-driven cancers require tbk1. *Nature*, 462(7269):108–112, 2009.
- JM Bernardo, MJ Bayarri, JO Berger, AP Dawid, D Heckerman, AFM Smith, and M West. Bayesian factor regression models in the “large p, small n” paradigm. *Bayesian statistics*, 7:733–742, 2003.
- Xin Bing, Florentina Bunea, Yang Ning, and Marten Wegkamp. Adaptive estimation in structured factor models with applications to overlapping clustering. *The Annals of Statistics*, 48(4):2055–2081, 2020.
- Eli Bingham, Jonathan P Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul Szerlip, Paul Horsfall, and Noah D Goodman. Pyro: Deep universal probabilistic programming. *The Journal of Machine Learning Research*, 20(1):973–978, 2019.
- Florian Buettner, Naruemon Pratanwanich, Davis J McCarthy, John C Marioni, and Oliver Stegle. f-sclvm: scalable and versatile factor analysis for single-cell rna-seq. *Genome biology*, 18(1):1–13, 2017.
- Yanshuai Cao and David J Fleet. Generalized product of experts for automatic and principled fusion of gaussian process predictions. *arXiv preprint arXiv:1410.7827*, 2014.
- Carlos M Carvalho, Nicholas G Polson, and James G Scott. Handling sparsity via the horseshoe. In *Artificial Intelligence and Statistics*, pages 73–80. PMLR, 2009.
- Carlos M Carvalho, Nicholas G Polson, and James G Scott. The horseshoe estimator for sparse signals. *Biometrika*, 97(2):465–480, 2010.
- AC Damianou, Carl Henrik Ek, MK Titsias, and ND Lawrence. Manifold relevance determination. In *29th International Conference on Machine Learning, ICML 2012, 26 June 2012 through 1 July 2012, Edinburgh*, pages 145–152, 2012.
- Stanislav Drápela, Prashant Khirsariya, Wytse M van Weerden, Radek Fedr, Tereza Suchánková, Diana Búzová, Jan Červený, Aleš Hampl, Martin Pühr, William R Watson, et al. The chk1 inhibitor mu380 significantly increases the sensitivity of human docetaxel-resistant prostate cancer cells to gemcitabine through the induction of mitotic catastrophe. *Molecular oncology*, 14(10):2487–2503, 2020.
- Barbara E Engelhardt and Matthew Stephens. Analysis of population structure: a unifying framework and novel methods based on sparse factor analysis. *PLoS genetics*, 6(9):e1001117, 2010.
- Antonio Fabregat, Steven Jupe, Lisa Matthews, Konstantinos Sidiropoulos, Marc Gillespie, Phani Garapati, Robin Haw, Bijay Jassal, Florian Korninger, Bruce May, et al. The reactome pathway knowledgebase. *Nucleic acids research*, 46(D1):D649–D655, 2018.
- H Robert Frost, Zhigang Li, and Jason H Moore. Principal component gene set enrichment (pcgse). *BioData mining*, 8(1):1–18, 2015.
- Kazutoshi Fujita and Norio Nonomura. Role of androgen receptor in prostate cancer: a review. *The world journal of men’s health*, 37(3):288–295, 2019.
- Soumya Ghosh, Jiayu Yao, and Finale Doshi-Velez. Structured variational learning of bayesian neural networks with horseshoe priors. In *International Conference on Machine Learning*, pages 1744–1753. PMLR, 2018.
- Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 2013.
- Harold Hotelling. Relations between two sets of variates. In *Breakthroughs in statistics*, pages 162–190. Springer, 1992.
- Minoru Kanehisa and Susumu Goto. Kegg: kyoto encyclopedia of genes and genomes. *Nucleic acids research*, 28(1):27–30, 2000.
- Karan Khanna, Nikki Salmond, Kalan S Lynn, Hon S Leong, and Karla C Williams. Clinical significance of steap1 extracellular vesicles in prostate cancer. *Prostate cancer and prostatic diseases*, 24(3):802–811, 2021.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

- Arto Klami, Seppo Virtanen, and Samuel Kaski. Bayesian canonical correlation analysis. *Journal of Machine Learning Research*, 14(4), 2013.
- Arto Klami, Seppo Virtanen, Eemeli Leppäaho, and Samuel Kaski. Group factor analysis. *IEEE transactions on neural networks and learning systems*, 26(9): 2136–2147, 2014.
- Hoon Young Kong and Jonghoe Byun. Emerging roles of human prostatic acid phosphatase. *Biomolecules & therapeutics*, 21(1):10, 2013.
- Andrew S Lan, Andrew E Waters, Christoph Studer, and Richard G Baraniuk. Sparse factor analysis for learning and content analytics. *The Journal of Machine Learning Research*, 15(1):1959–2008, 2014.
- Changhee Lee and Mihaela van der Schaar. A variational information bottleneck approach to multi-omics data integration. In *International Conference on Artificial Intelligence and Statistics*, pages 1513–1521. PMLR, 2021.
- Arthur Liberzon, Chet Birger, Helga Thorvaldsdóttir, Mahmoud Ghandi, Jill P Mesirov, and Pablo Tamayo. The molecular signatures database hallmark gene set collection. *Cell systems*, 1(6):417–425, 2015.
- Yuting Lin, Junichi Fukuchi, Richard A Hiipakka, John M Kokontis, and Jialing Xiang. Up-regulation of bcl-2 is required for the progression of prostate cancer cells from an androgen-dependent to an androgen-independent growth stage. *Cell research*, 17(6):531–536, 2007.
- Mohammad Lotfollahi, Sergei Rybakov, Karin Hrovatin, Soroor Hedyeh-Zadeh, Carlos Talavera-López, Alexander V Misharin, and Fabian J Theis. Biologically informed deep learning to query gene programs in single-cell atlases. *Nature Cell Biology*, pages 1–14, 2023.
- James Lucas, George Tucker, Roger B Grosse, and Mohammad Norouzi. Don’t blame the elbow! a linear vae perspective on posterior collapse. *Advances in Neural Information Processing Systems*, 32, 2019.
- David JC MacKay et al. Bayesian nonlinear modeling for the prediction competition. *ASHRAE transactions*, 100(2):1053–1062, 1994.
- Toby J Mitchell and John J Beauchamp. Bayesian variable selection in linear regression. *Journal of the american statistical association*, 83(404):1023–1032, 1988.
- Trevor Park and George Casella. The bayesian lasso. *Journal of the American Statistical Association*, 103(482): 681–686, 2008.
- Sven Perner, Niels J Rupp, Martin Braun, Mark A Rubin, Holger Moch, Manfred Dietel, Nicolas Wernert, Klaus Jung, Carsten Stephan, and Glen Kristiansen. Loss of slc45a3 protein (prostein) expression in prostate cancer is associated with slc45a3-erg gene rearrangement and an unfavorable clinical course. *International journal of cancer*, 132(4):807–812, 2013.
- Juho Piironen and Aki Vehtari. Sparsity information and regularization in the horseshoe and other shrinkage priors. *Electronic Journal of Statistics*, 11(2):5018–5051, 2017.
- Nicholas G Polson and James G Scott. Shrink globally, act locally: Sparse bayesian regularization and prediction. *Bayesian statistics*, 9(501-538):105, 2010.
- Rajesh Ranganath, Sean Gerrish, and David Blei. Black box variational inference. In *Artificial intelligence and statistics*, pages 814–822. PMLR, 2014.
- Veronika Ročková and Edward I George. The spike-and-slab lasso. *Journal of the American Statistical Association*, 113(521):431–444, 2018.
- Aravind Subramanian, Pablo Tamayo, Vamsi K Mootha, Sayan Mukherjee, Benjamin L Ebert, Michael A Gillette, Amanda Paulovich, Scott L Pomeroy, Todd R Golub, Eric S Lander, et al. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of the National Academy of Sciences*, 102(43):15545–15550, 2005.
- Valentine Svensson, Adam Gayoso, Nir Yosef, and Lior Pachter. Interpretable factor models of single-cell rna-seq via variational autoencoders. *Bioinformatics*, 36(11): 3418–3421, 2020.
- Louis Leon Thurstone. Multiple factor analysis. *Psychological review*, 38(5):406, 1931.
- Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- Katarzyna Tomczak, Patrycja Czerwińska, and Maciej Wiznerowicz. The cancer genome atlas (tcga): an immeasurable source of knowledge. *Contemporary oncology*, 19(1A):A68, 2015.
- Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- Shiwen Zhao, Chuan Gao, Sayan Mukherjee, and Barbara E Engelhardt. Bayesian group factor analysis with structured sparsity. *The Journal of Machine Learning Research*, 2016.

A APPENDIX A

Here, we provide additional details regarding the synthetic experiments. In particular, we extend the description of the baselines and their respective training settings. We then specify the evaluation metrics such as the RMSE for measuring the quality of the reconstruction, and additional binary metrics such as precision, recall and the F_1 score for assessing how well each model recovers the true underlying structure of the latent factors. We briefly describe the process of factor matching applied to all models in a post-processing step. Finally, we report additional results such as the precision and recall curves for the benchmark across all possible thresholds, and the performance of MuVI when dealing with significantly higher dimensional data.

A.1 Baselines

The parameter and training settings for the baselines are provided here. We train MOFA combining the ARD and the spike-and-slab prior on the weights for column-wise and element-wise sparsity. We apply BASS with default parameters and 20 parameter-expanded expectation maximization (PX-EM) iterations. The autoencoder consists of an encoder-decoder pair for each observed view, connected via a common latent space governed by a product of experts (PoE) approach (Lee and van der Schaar, 2021; Cao and Fleet, 2014). Each encoder comprises two hidden layers of size 64 and 32, followed by a ReLU activation. Each decoder then performs a linear mapping of the latent code to the observed features of each view. To achieve sparsity in the factor-to-feature mapping, we introduce an L1 penalty term with $\lambda = 0.01$ in the respective decoder weights. All models were trained until an early stopping condition was met, that is, until the optimization objective no longer improved in consecutive iterations, and after a certain number of patience steps.

A.2 Evaluation

To simplify notation, we omit the view index m and rewrite the collection of observations as a single matrix $\mathbf{Y} \in \mathbb{R}^{N \times D}$ comprising N samples across M coupled views $[\mathbf{Y}^{(1)}, \mathbf{Y}^{(2)}, \dots, \mathbf{Y}^{(M)}]$, where $\mathbf{Y}^{(m)} \in \mathbb{R}^{N \times D_m}$. Similarly, we represent the collection of the factor loadings $\mathbf{W} \in \mathbb{R}^{D \times K}$ as $[\mathbf{W}^{(1)}, \mathbf{W}^{(2)}, \dots, \mathbf{W}^{(M)}]$.

We assess the quality of each model by computing the reconstruction error measured by the RMSE,

$$RMSE(\mathbf{Y}, \hat{\mathbf{Y}}) = \sqrt{\frac{\sum_{i=1}^N \|\mathbf{y}_i - \hat{\mathbf{y}}_i\|^2}{N}},$$

where $\hat{\mathbf{y}}_i \in \mathbb{R}^D$ denotes the reconstructed sample \mathbf{y}_i . Next, we quantify the ability of each model to recover the true underlying structure of the factor loadings by computing the precision, recall and the F_1 score between the in-

Table 3: A confusion matrix between the inferred feature sets and the true feature sets.

| | actual active $ w_{j,k} > t$ | actual inactive $ w_{j,k} \leq t$ |
|--|----------------------------------|---------------------------------------|
| predicted active $ \hat{w}_{j,k} > t$ | TP | FP |
| predicted inactive $ \hat{w}_{j,k} \leq t$ | FN | TN |


ferred feature sets and the true feature sets. Here, the inferred feature sets refer to the binarized representation of the factor loadings $\hat{\mathbf{W}}$ in terms of active and inactive features based on a threshold, where factor loadings with an absolute value larger than the threshold are considered active, and, otherwise inactive. With a slight abuse in notation, let \mathbf{I}_k be the true feature set, and $\hat{\mathbf{I}}_k$ the inferred feature set for factor k , where $\hat{I}_{j,k} = 1$ if $|\hat{w}_{j,k}| > t$, and $\hat{I}_{j,k} = 0$ otherwise, for a given threshold t . Then, we may generate the confusion matrix in Table 3 to compute the precision, recall and the F_1 score.

$$Pr = \frac{TP}{TP + FP}, Rc = \frac{TP}{TP + FN}, F_1 = \frac{2 \cdot Pr \cdot Rc}{Pr + Rc}.$$

Finally, we describe the process of factor matching to overcome the issue of non-identifiability that is generally present in factor analysis. We rely on the cosine similarities between the inferred feature sets and the true feature sets to match the latent dimensions. Specifically, we compute a cosine similarity matrix $\mathbf{C} \in \mathbb{R}^{K \times K}$, where $c_{k,l} = \cos(\hat{\mathbf{I}}_k, \mathbf{I}_l)$, for $k, l \in \{1, \dots, K\}$, and compute an optimal permutation matrix $\mathbf{P} \in \mathbb{R}^{K \times K}$, such that the transformation $\mathbf{C}\mathbf{P}$ has the largest values in the diagonal, i.e. maximizes the trace $Tr(\mathbf{C}\mathbf{P})$. We then inspect the permutation matrix \mathbf{P} to match the factor dimensions between the inferred and the actual latent space.

A.3 Precision and Recall Curves

In addition to the threshold-based F_1 scores provided in 4.1.2, we show the precision and recall curves across all possible thresholds in Figure 5 for MuVI and the baseline models: BASS, MOFA and AE.

Figure 5: Threshold-agnostic precision and recall curves for $\text{MuVI}_{0.5}$ and the baseline models: BASS, MOFA, AE.Table 4: Performance evaluation when increasing the number of features D_m for each of the four views m , while keeping the number of samples low, $N = 200$. Average scores across five independent runs \pm standard deviation.

| | RMSE | Precision | Recall | F_1 |
|-------------|------------------|------------------|------------------|------------------|
| $D_m = 1K$ | 0.322 ± 0.00 | 0.937 ± 0.03 | 0.995 ± 0.00 | 0.965 ± 0.02 |
| $D_m = 5K$ | 0.323 ± 0.00 | 0.943 ± 0.02 | 0.995 ± 0.00 | 0.968 ± 0.01 |
| $D_m = 10K$ | 0.323 ± 0.00 | 0.948 ± 0.02 | 0.996 ± 0.00 | 0.971 ± 0.00 |

A.4 Performance in Higher Dimensions

We repeat the synthetic experiments by exacerbating the gap between the number of samples and the number of features ($N \ll D$). For each of the four views m , we increase the corresponding number of features D_m to one, five and ten thousand features, while keeping the number of samples low, $N = 200$. The results reported in Table 4 for the uninformed version of our model (MuVI_U) are consistent with the results shown in 4.1.2.

B APPENDIX B

Here, we further investigate the interplay between the quality of the prior information, and the degree of the prior belief encoded by the single hyperparameter α .

B.1 Sensitivity to the Prior Penalty

Figure 8 summarizes the ablation study for the choice of α across different noise levels, in terms of the recall (left) and precision (right) between the inferred and the true feature sets, when informing only a single view (view 0). A larger α (top) translates to a weaker penalty to the features not present in the (noisy) feature set, whereas a smaller alpha (bottom) encodes a stronger prior belief, and therefore a larger penalty. As the prior belief increases, i.e. α

Table 5: Performance comparison across all noise levels and penalties in the ablation study. Average RMSE across five independent runs \pm standard deviation.

| | $\alpha = 0.1$ | $\alpha = 0.05$ | $\alpha = 0.03$ | $\alpha = 0.01$ |
|---------------------|------------------|------------------|------------------|------------------|
| $\text{MuVI}_{0.0}$ | 0.321 ± 0.00 | 0.321 ± 0.00 | 0.322 ± 0.00 | 0.328 ± 0.02 |
| $\text{MuVI}_{0.1}$ | 0.321 ± 0.00 | 0.321 ± 0.00 | 0.322 ± 0.00 | 0.331 ± 0.03 |
| $\text{MuVI}_{0.2}$ | 0.321 ± 0.00 | 0.321 ± 0.00 | 0.322 ± 0.00 | 0.335 ± 0.02 |
| $\text{MuVI}_{0.5}$ | 0.321 ± 0.00 | 0.321 ± 0.00 | 0.322 ± 0.00 | 0.339 ± 0.03 |
| $\text{MuVI}_{0.9}$ | 0.321 ± 0.00 | 0.322 ± 0.00 | 0.326 ± 0.01 | 0.367 ± 0.03 |
| $\text{MuVI}_{1.0}$ | 0.321 ± 0.00 | 0.322 ± 0.00 | 0.328 ± 0.01 | 0.393 ± 0.03 |

decreases, we observe an increase in the amount of false negatives as the noise level increases. This behavior is to be expected, since the noisy prior incorrectly penalizes the majority of the true signals. In the bottom row, in particular, the high penalty induced for the uninformed views, i.e. view 1, 2 and 3, reduces their learning capacity relative to the informed view, which exhibits no difficulties in inferring the true signals. At the same time, the amount of false positives decreases, resulting in a higher precision score when having a higher degree of belief in the prior information. In other words, a weaker belief may lead to additional and redundant signals escaping the regularized horseshoe penalty, and a stronger belief may restrict some of the true signals, depending on the amount of noise in the prior information. The ablation study indicates an $\alpha \approx 0.03$ performs the best in terms of both precision and recall, and strikes a better balance between recovering the true positive signals while restricting the true negative signals. Nevertheless, the quality of the latent representation is preserved across all noise levels and penalties. In Table 5 we report the reconstruction error for all training configurations involved in the ablation study. The RMSE obtained for the majority of the models still outperforms the baselines, and only drops slightly when combining highly perturbed prior feature sets with a high penalty.

B.2 Learning Identifiable Factors from Pure Variables

In Section 4.1.4, in Figure 3 we assess to what extent our model preserves the identifiability properties when increasing the amount of noise in the prior annotations. The noise level serves as a proxy for the number of pure variables in the prior information, that is, a higher noise level results in fewer pure variables a priori. A noise fraction of 100%, for instance, results in zero pure variables, and therefore makes the model unidentifiable. In Figure 6 we quantify the number of known pure variables corresponding to each noise fraction. The results are aggregated across all feature sets generated for the synthetic datasets (see 4).




Figure 6: Noise fraction as a proxy for the number of pure variables. Each boxplot summarizes the distribution of the pure variables across all factors for the feature sets generated in the synthetic experiments, for different noise levels.

Next, we investigate the minimal number of pure variables required in a feature set to render the corresponding latent factor identifiable. In Figure 7 we count the number of the pure variables and compute the F_1 score for each factor individually, without performing factor matching. In other words, we compare the inferred feature set $\hat{\mathbf{I}}_k$ with the true feature set \mathbf{I}_k for the same latent dimension k . Since several factors may have the same number of pure variables, we report the median scores and a 95% confidence interval. The results strongly suggest that as few as two pure variables are sufficient to guarantee identifiable factors, in accordance with the theory on the identifiability of factor analysis (Anderson and Rubin, 1956).




Figure 7: Inferring identifiable factors from pure variables. The median F_1 score when increasing the number of pure variables available in the prior feature set, shaded by the 95% confidence interval. The inferred and true factors are compared along the same dimension, without performing factor matching.

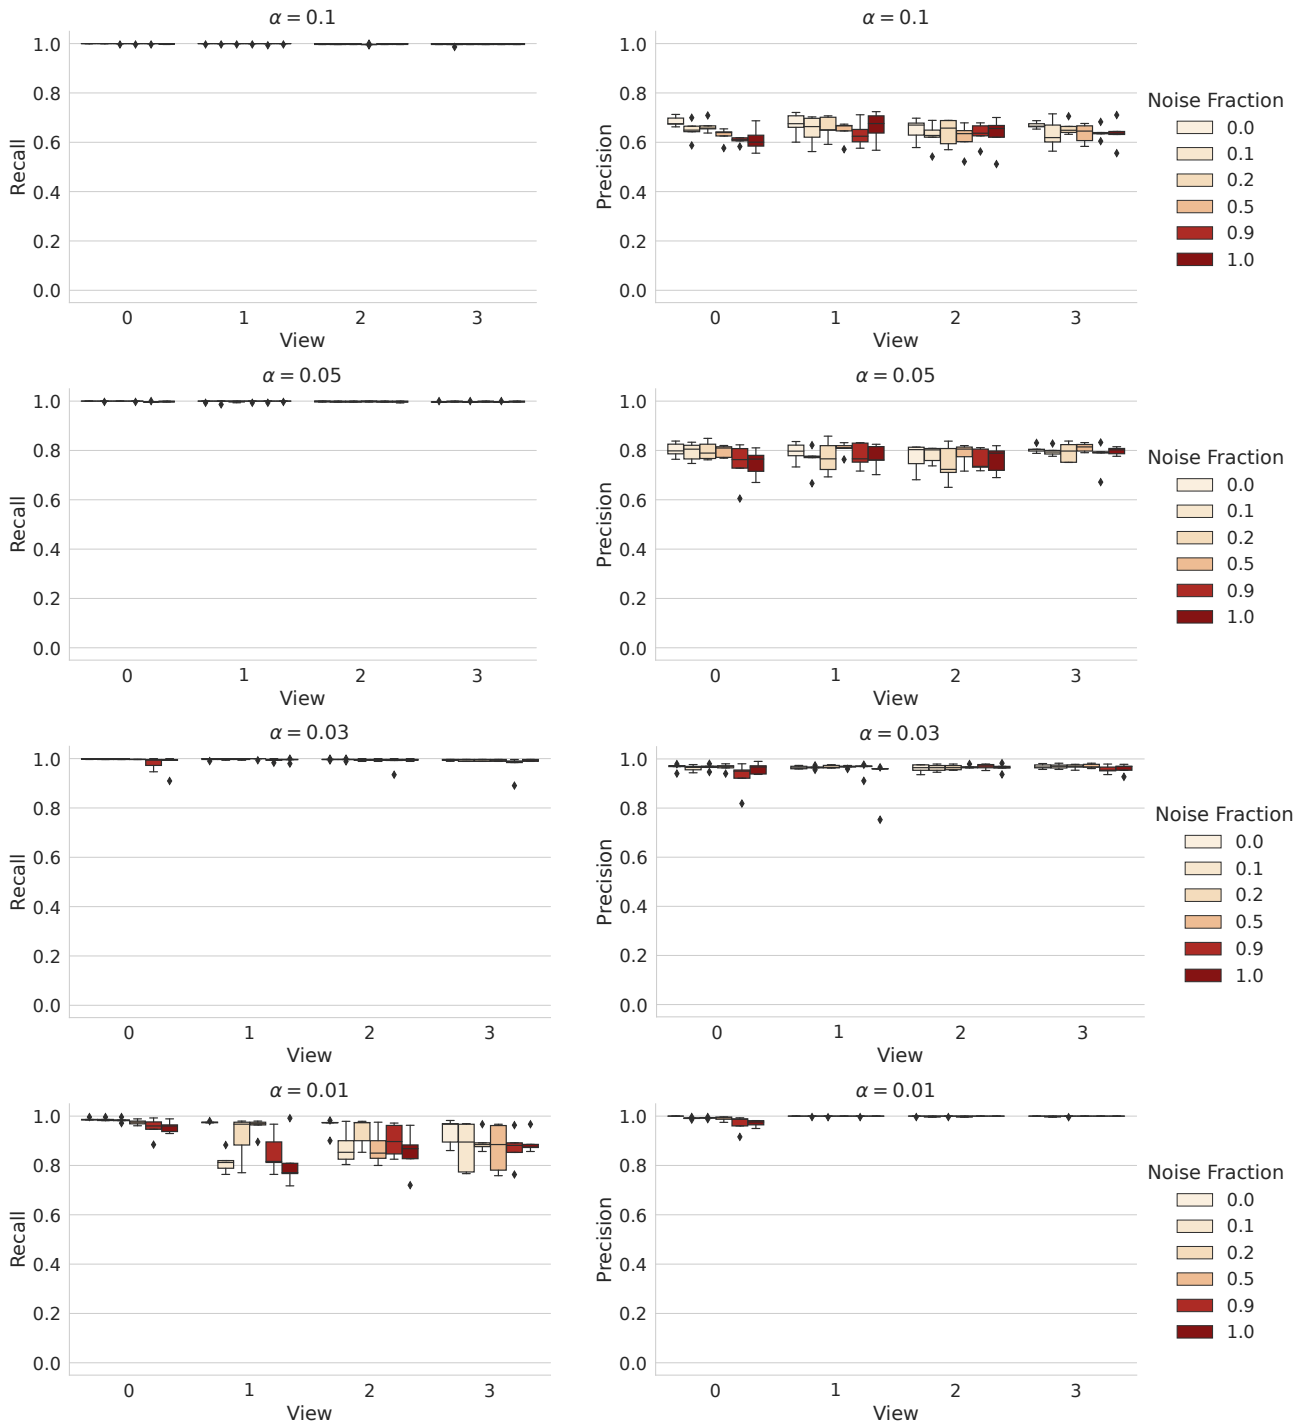


Figure 8: Sensitivity analysis for the choice of α across different noise levels, in terms of the recall (left) and precision (right) between the inferred and the true feature sets. A larger α (top) translates to a weaker penalty to the features not present in the (noisy) feature set, whereas a smaller alpha (bottom) encodes a stronger prior belief, and therefore a larger penalty. Results when informing only the first view.

C APPENDIX C

We complement the analysis on the real TCGA dataset from Section 4.2 with additional results.

C.1 Availability of Pure Variables in the Gene Sets

In Figure 9, we provide a summary of the number of pure variables available in each feature set (gene set) collection. Due to being more specialized, Hallmark and KEGG include more pure variables than the more general collection of Reactome. However, we may only pinpoint the pure variables according to the definition of the gene sets. In practice, we make no assumptions about the gene sets or the number of pure variables available in each gene set, and treat them as partially correct prior information, thereby allowing the model to refine each pathway based on the training data.




Figure 9: Number of pure variables in each gene set (feature set) collection integrated in the TCGA analysis.

In order to determine whether the inferred factor loadings differ significantly from the features present in the prior gene set annotation, we perform an adjusted parametric t-test (Frost et al., 2015) based on the correlation between the features that are present in the prior annotation.

C.2 Latent Space Inferred by the Baselines

For a visual comparison, we apply a t-SNE approach to the latent representations inferred by each baseline model and map the data onto the two dimensional embeddings as shown in Figure 10.

C.3 Matching Relevant Factors to Cancer Types

Finally, we provide a summary of the relevant pathways for each cancer subtype. We perform a one-vs-rest Wilcoxon rank test to identify the main pathways that capture the differences among patients of different cancer types. In Figure 11, we report the top 3 pathways corresponding to each cancer subtype. The test results indicate that both the androgen response and the melanogenesis pathway are highly relevant for prostate cancer (PRAD) and melanoma related cancer subtypes (SKCM and UVM). In addition, we apply a standard hierarchical clustering on the factor scores to group similar cluster subtypes together. Among others, we observe biologically meaningful structures such as identifying similarities across the only two groups of melanoma samples such as the skin cutaneous melanoma (SKCM) and uveal melanoma (UVM).




Figure 10: Embedding space learned by the baselines: BASS, MOFA and the hybrid autoencoder (AE).

Encoding Domain Knowledge in Multi-view Latent Variable Models



Figure 11: Relevant pathways for each cancer subtype. A one-vs-rest Wilcoxon rank test identifies the main pathways that drive the heterogeneity of patients with respect to their cancer types. Each group of the top 3 pathways corresponds to a single cancer subtype on the right. The color intensity for each cell depicts the average factor score in each patient subpopulation. Cluster subtypes ordered via a standard hierarchical clustering on their factor scores.