

# A BIAS-VARIANCE-COVARIANCE DECOMPOSITION OF KERNEL SCORES FOR GENERATIVE MODELS

Anonymous authors

Paper under double-blind review

## ABSTRACT

Generative models, like large language models, are becoming increasingly relevant in our daily lives, yet a theoretical framework to assess their generalization behavior and uncertainty does not exist. Particularly, the problem of uncertainty estimation is commonly solved in an ad-hoc manner and task dependent. For example, natural language approaches cannot be transferred to image generation. In this paper we introduce the first bias-variance-covariance decomposition for kernel scores and their associated entropy. We propose unbiased and consistent estimators for each quantity which only require generated samples but not the underlying model itself. As an application, we offer a generalization evaluation of diffusion models and discover **evidence of** how mode collapse is a contrary phenomenon to overfitting. Further, we demonstrate that variance and predictive kernel entropy are viable measures of uncertainty for image, audio, and language generation. Specifically, our approach for uncertainty estimation is more predictive of performance on CoQA and TriviaQA question answering datasets than existing baselines and can also be applied to closed-source models.

## 1 INTRODUCTION

In recent years, generative models have revolutionized daily lives well beyond the field of machine learning (Kasneci et al., 2023; Meskó & Topol, 2023). These models have found applications in diverse domains, including image creation (Ramesh et al., 2021), natural language generation (OpenAI, 2023), drug discovery (Paul et al., 2021), and speech synthesis (Ning et al., 2019). While generative models have demonstrated remarkable capabilities in generating data that closely resemble real-world samples, they often fall short in providing the vital and often overlooked aspect of uncertainty estimation (Wu & Shang, 2020). Uncertainty estimation in machine learning is a critical component of model performance assessment and deployment (Hekler et al., 2023). It addresses the inherent limitations and challenges associated with machine learning based decisions. For generative models, this may include their propensity to generate improbable or nonsensical samples (“hallucinations”). Even though uncertainty estimation methods for natural language question

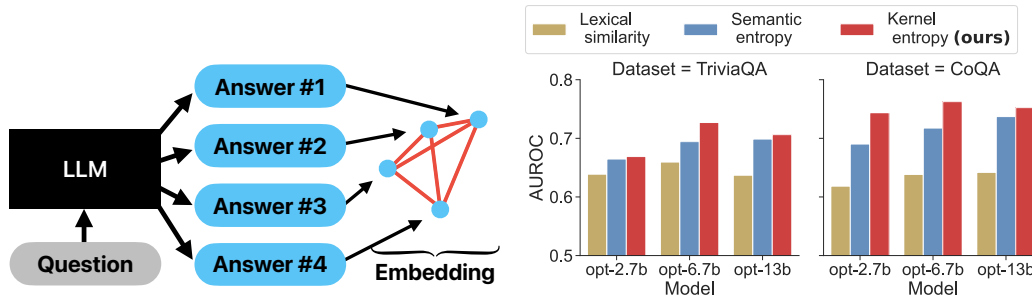


Figure 1: **Left:** Illustration of predictive kernel entropy for a generative model. A kernel measures the pairwise similarities (red lines) of outputs in a vector space. The predictive kernel entropy is then the negative average kernel value. **Right:** The predictive kernel entropy shows the best performance among uncertainty approaches for closed-source settings (c.f. Section 5.3).

answering tasks exist (Kuhn et al., 2023), they are ad-hoc without theoretical grounding and are not transferable to other data generation tasks .

Predictive uncertainty is an informal concept, but it is implied that it relates to the prediction error without requiring access to target outcomes. A formal approach to this is the bias-variance decomposition, a central concept in statistical learning theory (Bishop & Nasrabadi, 2006; Hastie et al., 2009; Murphy, 2022). It helps to understand the generalization behavior of models and naturally raises uncertainty terms by isolating the target prediction into a bias term (Gruber & Buettner, 2023). Ueda & Nakano (1996) discovered the bias-variance-covariance decomposition of the mean squared error, which is the foundation of negative correlation learning (Liu & Yao, 1999b;a; Brown, 2004) and for reducing correlations in weight averaging (Rame et al., 2022). Though the bias-variance decomposition has been generalized to distributions (Gruber & Buettner, 2023), the current theory does not include a covariance term and relies on having access to the predicted distribution. But, many generative models only indirectly fit the training distribution by learning how to generate samples. Others, such as large language models (LLMs), do explicitly fit the target distribution, but the prevalence of closed source models means the predictive distribution is often not available to the practitioner (OpenAI, 2023). This makes it infeasible to apply the powerful framework of the bias-variance decomposition in these cases.

Contrary, kernels allow to quantify differences between distributions only based on their samples without requiring access to these distributions (Gretton et al., 2012a). They are used in kernel scores to assess the goodness-of-fit for predicted distributions (Gneiting & Raftery, 2007).

As **contribution** in this work, we...

- introduce the first extension of the bias-variance-covariance decomposition beyond the mean squared error to kernel scores in Section 3, and propose unbiased and consistent estimators only requiring generated samples in Section 4.
- examine the generalisation behavior of generative models for image and audio generation and investigate how bias, variance and kernel entropy relate to the generalisation error in Section 5. This includes **evidence** that mode collapse of underrepresented minority groups is expressed purely in the bias.
- demonstrate how kernel entropy in combination with text embeddings outperforms existing methods for estimating the uncertainty of LLMs on common question answering datasets (c.f. Figure 1 and Section 5.3).

## 2 BACKGROUND

In this section, we give a brief introduction into kernel scores, followed up by other bias-variance decompositions and approaches for assessing the uncertainty in natural language generation.

### 2.1 KERNEL SCORES

Kernel scores are a class of loss functions for distribution predictions (Eaton, 1981; Eaton et al., 1996; Dawid, 2007). For simplicity, we omit complex-valued kernels. We refer to a symmetric kernel  $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$  defined on a set  $\mathcal{X}$  as **positive definite** (p.d.) if  $\sum_{i=1}^n \sum_{j=1}^n a_i k(x_i, x_j) a_j > 0$  for all  $x_1, \dots, x_n \in \mathcal{X}$  and  $a_1, \dots, a_n \neq 0$  with  $n \in \mathbb{N}$ . **Positive semi-definite** (p.s.d.) refers to the case when only ' $\geq$ ' holds. Assume  $\mathcal{P}$  is a set of distributions defined on  $\mathcal{X}$  such that for a kernel  $k$  the operator  $\langle P | k | Q \rangle := \int_{\mathcal{X}} \int_{\mathcal{X}} k(x, y) dP(x) dQ(y)$  is finite for all  $P, Q \in \mathcal{P}$  (Eaton, 1981). It follows that  $\langle \cdot | k | \cdot \rangle$  is a symmetric bilinear form and induces the semi-norm  $\|P\|_k = \sqrt{\langle P | k | P \rangle}$ . A **kernel score**  $S_k: \mathcal{P} \times \mathcal{X} \rightarrow \mathbb{R}$  based on a p.s.d. kernel  $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$  is defined as (Steinwart & Ziegel, 2021)

$$S_k(P, y) = \|P\|_k^2 - 2 \langle P | k | \delta_y \rangle, \quad (1)$$

where  $\delta_y$  is the dirac measure at point  $y$ . Note that Eaton (1981) and Dawid (2007) use a slightly less general definition. If  $\mathcal{P}$  only consists of Borel probability measures, then the expected kernel score  $\mathbb{E}[S_k(P, Y)]$  based on a target  $Y \sim Q \in \mathcal{P}$  is minimized when  $P = Q$  (Gneiting & Raftery, 2007). Following Dawid (2007), we refer to  $-\|Q\|_k^2 = \mathbb{E}[S_k(Q, Y)]$  as the **kernel entropy** function of  $Q$ . If  $k$  is associated with a reproducing kernel Hilbert space (RKHS), then the kernel score is connected to the **maximum mean discrepancy** (MMD) via  $\text{MMD}_k^2(P, Q) = \mathbb{E}[S_k(P, Y)] + \|Q\|_k^2$

(Steinwart & Ziegel, 2021). MMDs are used for non-parametric two-sample testing (Gretton et al., 2012a) and generative image modelling (Li et al., 2015; Bińkowski et al., 2018). Compared to MMDs, kernel scores are applicable to a wider range of scenarios, since one sample of the target distribution is sufficient for evaluation. For example, MMDs cannot be computed for question-answering pairs when there is only one answer for each question in the dataset.

## 2.2 BIAS-VARIANCE (-COVARIANCE) DECOMPOSITIONS

Ueda & Nakano (1996) introduced the bias-variance-covariance decomposition for the mean squared error. For a real-valued ensemble prediction  $\hat{P}^{(n)} = \frac{1}{n} \sum_{i=1}^n \hat{P}_i$  with identically distributed  $\hat{P}_1, \dots, \hat{P}_n$  and real-valued target  $Y$  it is given by

$$\underbrace{\mathbb{E} \left[ \left( \hat{P}^{(n)} - Y \right)^2 \right]}_{\text{Expected Squared Error}} = \underbrace{\mathbb{V}(Y)}_{\text{Noise}} + \underbrace{\left( \mathbb{E}[\hat{P}] - \mathbb{E}[Y] \right)^2}_{\text{Bias}} + \underbrace{\frac{1}{n} \mathbb{V}(\hat{P})}_{\text{Variance}} + \underbrace{\frac{n-1}{n} \text{Cov}(\hat{P}, \hat{P}')}_{\text{Covariance}}, \quad (2)$$

with  $\hat{P} := \hat{P}_1$  and  $\hat{P}' := \hat{P}_2$ . Rame et al. (2022) propose an approximate bias-variance-covariance decomposition for hard-label classification but it only holds in an infinitesimal locality around the prediction. To our best knowledge, the mean squared error is the only case so far with a non-approximated decomposition. Gruber & Buettner (2023) introduced a bias-variance decomposition for loss functions of general distributions. They demonstrated that the variance term is a meaningful measure of the model uncertainty similar to confidence scores in classification. But, their formulation requires a loss-specific transformation of the distributions into a dual vector space and a covariance term is not given.

## 2.3 UNCERTAINTY IN NATURAL LANGUAGE GENERATION

In the following, we give a brief overview of uncertainty estimations in natural language generation. A common approach is predictive entropy, which is the Shannon entropy  $-\int \log \hat{p}(y | x) d\hat{p}(y | x)$  of the predicted distribution  $\hat{p}$  given an input  $x$  (Malinin & Gales, 2020). For a generated token sequence  $\mathbf{s} = (s_1, \dots, s_l) \in \mathbb{N}^l$  of length  $l \in \mathbb{N}$  it is computed via  $\sum_{i=1}^l \log \hat{p}(s_i | s_1, \dots, s_{i-1})$ , where  $\hat{p}$  is the predicted distribution of the generating language model. Note that the predicted distribution is not always available for closed-source models. The computation also scales linearly with the length of the generated text, making it costly for larger text generations. Malinin & Gales (2020) propose to use length-normalisation of the predictive entropy since the Shannon entropy is systematically affected by the sequence length. Kuhn et al. (2023) propose *semantic entropy* to ease the computation of the predictive entropy by finding clusters of semantically similar generations. Another approach is *lexical similarity* (Fomicheva et al., 2020), which quantifies the average pairwise similarity between generated answers according to a similarity measure, like RougeL (Lin & Och, 2004; Kuhn et al., 2023). Kadavath et al. (2022) propose the baseline  $p(\text{True})$ , which asks the model itself if the generated answer is correct. Alternative approaches exist, which require an ensemble of models (Lakshminarayanan et al., 2017; Malinin & Gales, 2020). However, ensembles are practically less relevant due to the high computational cost of training even a single model.

## 3 A BIAS-VARIANCE-COVARIANCE DECOMPOSITION OF KERNEL SCORES

In this section, we state our main theoretical contribution. All proofs are presented in Appendix D. To highlight the similarity to the mean squared error case, we introduce the novel definitions for distributional variance and distributional covariance. The latter also implies a distributional correlation, which we define later in Section 4. Note that conventional variance and covariance are based on multiplication of two components ( $x \cdot x$  for variance and  $x \cdot y$  for covariance). We interpret  $\langle \cdot | k | \cdot \rangle$  as a generalization of this multiplication, which directly implies the following.

**Definition 3.1.** Assume we have a p.s.d. kernel  $k$  and random variables  $P$  and  $Q$  with outcomes in a distribution space as defined above. We define the **distributional variance** generated by  $k$  of  $P$  as

$$\text{Var}_k [P] = \mathbb{E} [\|P - \mathbb{E}[P]\|_k^2] \quad (3)$$

and the **distributional covariance** generated by  $k$  between  $P$  and  $Q$  as

$$\text{Cov}_k (P, Q) = \mathbb{E} [\langle P - \mathbb{E}[P] | k | Q - \mathbb{E}[Q] \rangle]. \quad (4)$$

If  $P$  is deterministic, i.e. is a random variable with only one outcome, then  $\text{Var}_k(P) = 0$ . Further, we have  $\text{Cov}_k(P, P) = \text{Var}_k[P]$ , and, if  $P$  and  $Q$  are independent, then  $\text{Cov}_k(P, Q) = 0$ . Note that the terms *kernel variance* and *kernel covariance* already exist in the literature and should not be confused with our definitions (Gretton et al., 2003).

We now have the necessary tools to state our main theoretical contribution in a concise manner.

**Theorem 3.2.** *Let  $S_k$  be a kernel score based on a p.s.d. kernel  $k$  and  $\hat{P}$  a predicted distribution for a target  $Y \sim Q$ , then*

$$\underbrace{\mathbb{E}[S_k(\hat{P}, Y)]}_{\text{Generalization Error}} = \underbrace{-\|Q\|_k^2}_{\text{Noise}} + \underbrace{\|\mathbb{E}[\hat{P}] - Q\|_k^2}_{\text{Bias}} + \underbrace{\text{Var}_k(\hat{P})}_{\text{Variance}}. \quad (5)$$

If we have an ensemble prediction  $\hat{P}^{(n)} := \frac{1}{n} \sum_{i=1}^n \hat{P}_i$  with identically distributed members  $\hat{P}_1, \dots, \hat{P}_n$ , then

$$\text{Var}_k(\hat{P}^{(n)}) = \frac{1}{n} \text{Var}_k(\hat{P}_1) + \frac{n-1}{n} \text{Cov}_k(\hat{P}_1, \hat{P}_2). \quad (6)$$

This theorem proves the relation between a kernel-based generalization error and the distributional variance and distributional covariance. It has a wide range of practical relevance since kernels can be used for almost all data scenarios via vector embeddings (Liu et al., 2020). Consequently, we can extend the evaluation and analysis for regression, which has been done since the emergence of the mean squared error decomposition (Brown, 2004), to tasks with arbitrarily complex target distributions. This opens up possibilities for gaining new insights into the most successful generative models of recent years. **In Appendix B we express the decomposition in terms of the reproducing kernel Hilbert space associated with the kernel.**

In the following of this work, we use Theorem 3.2 to study the generalization behavior of generative models and to find ways to estimate the uncertainty of generated data. The presented evaluations and approaches are applicable to almost any data generation task due to the flexibility of kernels and data embeddings.

**Predictive Kernel Entropy for Single Models.** Historically, the bias-variance decomposition had a large impact on the development of some of the most established machine learning algorithms, like Random Forests (Breiman, 2001) or Gradient Boosting (Friedman, 2002). However, ensemble approaches are not similarly dominant for generative modeling. Estimating  $\text{Var}_k(\hat{P})$  requires an ensemble of models, which is not always feasible. Instead, note the decomposition  $\text{Var}_k(\hat{P}) = \mathbb{E}[\|\hat{P}\|_k^2] - \|\mathbb{E}[\hat{P}]\|_k^2$  and observe that the distributional variance depends on the predictive kernel entropy  $-\|\hat{P}\|_k^2$ , which is estimated for single models. The predictive kernel entropy also appears in the definition of kernel scores in Equation 1. This suggests that it may have a substantial influence on the generalization error. In Section 5, we will discover that this influence is extremely high (Pearson correlation of approx. 0.95), but the sign of the correlation is task-specific. Further, by using text embeddings, predictive kernel entropy is better than other baselines in predicting the performance of LLMs (c.f. Section 5).

## 4 UNBIASED AND CONSISTENT ESTIMATORS

If the prediction  $\hat{P}$  is available in closed-form, the quantities in Theorem 3.2 can be computed according to conventional approaches (Gruber & Buettner, 2023). But, this is not the case for a lot of recently developed generative models in Deep Learning. For example, Diffusion Models (Ho et al., 2020) or closed-source LLMs (OpenAI, 2023) are also learning the training distribution, but they are often limited to generating samples. In this section, we introduce estimators of the distributional variance and covariance for the case when only samples of the distributions are available. This increases the practical applicability of Theorem 3.2 by a wide margin and allows investigating the most recent and largest generative models without constraints. We assume a minimum of two samples from each distribution is given. All estimators in the following require a two-stage sampling procedure (Särndal et al., 2003): First, distributions are sampled in an outer loop, which can be seen as clusters. In Section 5, this will be an ensemble of generative models. Second, we sample of each

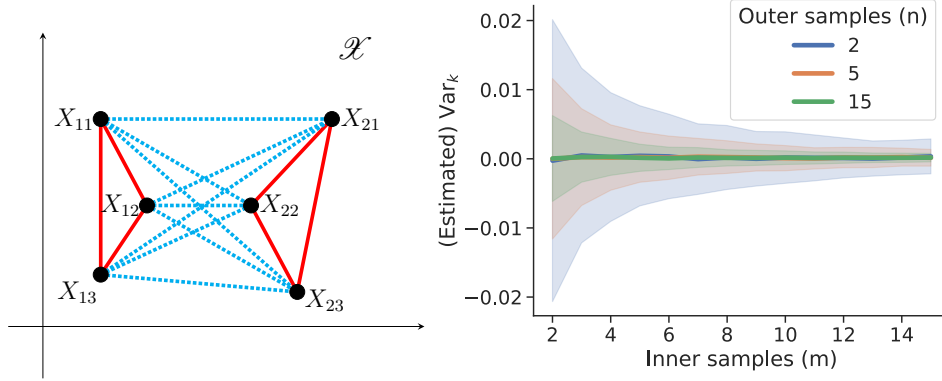


Figure 2: **Left:** Illustration of the estimator  $\widehat{\text{Var}}_k^{(n,m)}$  in the sample space  $\mathcal{X}$  for  $n = 2$  outer samples and  $m = 3$  inner samples. The estimator computes the average similarity within clusters (solid red lines) minus the average similarity between clusters (dotted blue lines). Shorter lines indicate higher similarity and larger kernel values. **Right:** Estimator standard deviation for various sample sizes. Even though the estimator does not converge in theory with the inner sample size  $m$ , it may still be influenced significantly by it for small sample sizes.

distribution multiple times in an inner loop, which can be seen as within-cluster samples. This will be the data generations of each model.

The procedure differs slightly between the variance and covariance case. For simplicity, we also assume that all within-cluster sample sizes are the same. All estimators can be adjusted if that is not the case and will still be unbiased and consistent. Again, all proofs are presented in Appendix D.

#### 4.1 DISTRIBUTIONAL VARIANCE

Assume we have a random variable  $P$  with outcomes in  $\mathcal{P}$  based on an unknown distribution  $\mathbb{P}_P$  from which we can sample. First, we sample distributions  $P_1, \dots, P_n \stackrel{\text{iid}}{\sim} \mathbb{P}_P$ . Then, we sample  $X_{i1}, \dots, X_{im} \stackrel{\text{iid}}{\sim} P_i$  for  $i = 1 \dots n$ . The estimator we are about to propose is directly derived from the conventional variance estimator  $\hat{\sigma}^2 := \frac{1}{n-1} \sum_{i=1}^n \|P_i - \frac{1}{n} \sum_{s=1}^n P_s\|_k^2$ . Note that it holds  $\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n \|P_i\|_k^2 - \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{s=1, s \neq i}^n \langle P_i | k | P_s \rangle$ , i.e. the estimator is the average of same-index pairs minus the average of the rest. Our extended estimator then uses the plug-ins  $\|P_i\|_k^2 \approx \frac{1}{m(m-1)} \sum_{j=1}^m \sum_{t=1, t \neq j}^m k(X_{ij}, X_{it})$  and  $\langle P_i | k | P_s \rangle \approx \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, X_{st})$ . The complete estimator of the distributional variance  $\text{Var}_k(P)$  is defined by

$$\widehat{\text{Var}}_k^{(n,m)} = \underbrace{\frac{1}{nm(m-1)} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m k(X_{ij}, X_{it})}_{\text{Average similarity within clusters}} - \underbrace{\frac{1}{n(n-1)m^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, X_{st})}_{\text{Average similarity between clusters}}. \quad (7)$$

An illustration is given on the left in Figure 2. The estimator is unbiased since  $\mathbb{E}[\widehat{\text{Var}}_k^{(n,m)}] = \text{Var}_k(P)$ . Its runtime complexity is in  $\mathcal{O}(m^2n^2)$ . Estimators with lower complexity, like  $\mathcal{O}(mn)$ , exist but are not recommendable since they have a worse performance and in most applications, generating the samples is far more costly than evaluating the estimator.

The variance of the estimator is in  $\mathcal{O}(\frac{1}{n}(1 + \frac{1}{m}))$ , which proves  $\widehat{\text{Var}}_k^{(n,m)} \rightarrow \text{Var}_k(P)$  in probability with growing  $n$  but not  $m$ . In words, the estimator is consistent with increasing outer samples but not inner samples. This may suggest to neglect creating inner samples and keep  $m$  small, but our analysis in Appendix D.2 shows that there exist sub-terms which converge equally fast in  $m$  as in  $n$ . In combination with the finite sample simulation in Figure 2, we recommend to use  $n \approx m \geq 10$ , if no prior information is available.

## 4.2 DISTRIBUTIONAL COVARIANCE AND CORRELATION

For the covariance case, assume we have random variables  $P$  and  $Q$  with outcomes in  $\mathcal{P}$  based on an unknown joint distribution  $\mathbb{P}_{PQ}$  from which we can sample. We require samples  $X_{i1}, \dots, X_{im} \stackrel{\text{iid}}{\sim} P_i$  and  $Y_{i1}, \dots, Y_{im} \stackrel{\text{iid}}{\sim} Q_i$  with  $(P_1, Q_1), \dots, (P_n, Q_n) \stackrel{\text{iid}}{\sim} \mathbb{P}_{PQ}$ . Then, we propose the unbiased and consistent covariance estimator

$$\widehat{\text{Cov}}_k^{(n,m)}(\mathbf{X}, \mathbf{Y}) = \frac{1}{nm^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m \left( k(X_{ij}, Y_{it}) - \frac{1}{n-1} \sum_{\substack{s=1 \\ s \neq i}}^n k(X_{ij}, Y_{st}) \right) \quad (8)$$

with  $\mathbf{X} := (X_{ij})_{i=1 \dots n, j=1 \dots m}$  and  $\mathbf{Y} := (Y_{ij})_{i=1 \dots n, j=1 \dots m}$ . It has the same runtime complexity and convergence rate as the variance estimator of Equation 7 (c.f. Appendix D.3). While the distributional covariance is directly implied by Theorem 3.2, it is difficult to interpret since it is not bounded. Consequently, we propose the distributional correlation estimator based on Equation 8 given by

$$\widehat{\text{Corr}}_k^{(n,m)} = \frac{\widehat{\text{Cov}}_k^{(n,m)}(\mathbf{X}, \mathbf{Y})}{\sqrt{\widehat{\text{Cov}}_k^{(n,m)}(\mathbf{X}, \mathbf{X})} \sqrt{\widehat{\text{Cov}}_k^{(n,m)}(\mathbf{Y}, \mathbf{Y})}} \in [-1, 1]. \quad (9)$$

It is consistent since continuous transformations of consistent estimators are also consistent (Shao, 2003), i.e. for  $n \rightarrow \infty$  in probability

$$\widehat{\text{Corr}}_k^{(n,m)} \rightarrow \text{Corr}_k(P, Q) := \frac{\text{Cov}_k(P, Q)}{\sqrt{\text{Var}_k(P)} \sqrt{\text{Var}_k(Q)}}. \quad (10)$$

We use the covariance estimator for the variance terms since the variance estimator can be negative and the correlation estimator is only asymptotically unbiased no matter the choice.

In the next section, we show that distributional correlation, as implied by Theorem 3.2, is a natural tool to gain new insights into the fitting process of generative models. For example, the correlations between epochs indicate how stable the convergence during training is.

## 5 APPLICATIONS

In this section, we apply the proposed statistical tools to assess generative models across a variety of different data generation tasks. **Specifically, we put emphasis on instance-level uncertainty estimation. A meaningful measure of uncertainty is able to predict the loss for a given prediction. We use the Pearson correlation coefficient to quantify how well an uncertainty measure predicts a continuous loss, and the area under receiver operator characteristic (AUROC) for a binary loss.** We first start in Section 5.1 with diffusion models for image generation on the synthetic InfiMNIST dataset for a detailed examination of their generalization behavior. In Section 5.2, we evaluate an ensemble of Glow-TTS models for text-to-speech synthesis on the SpeechLJ dataset. **In all cases, kernel entropy shows strong performance as uncertainty measure.** Last, we use kernel entropy to outperform other baselines in uncertainty estimation for natural language generation in Section 5.3. There, we evaluate single OPT models of different sizes on the question answering datasets CoQA and TriviaQA. The source code of all experiments is openly available at [https://github.com/\[waiting-for-acceptance\]](https://github.com/[waiting-for-acceptance]).

### 5.1 IMAGE GENERATION

For image generation, we use conditional diffusion models (Ho et al., 2020; Ho & Salimans, 2021) trained on MNIST-like datasets. We use InfiMNIST to sample an infinite number of uniquely perturbed MNIST images (Loosli et al., 2007). By simulating the data generation process we can assess the ground truth generalization error of the diffusion model. We sample  $n = 20$  distinct training sets from InfiMNIST each of size 60.000 (similar as MNIST). We then train a model on each training set. This is in correspondence to how generalization error, bias, and variance are evaluated for regression and classification tasks (Ueda & Nakano, 1996; Gruber & Buettner, 2023). We use  $m = 20$  generated images per class and per model for all estimators. The predictive kernel entropy is only evaluated on a single model to stay as closely as possible to practical constraints. Our kernel choice is the commonly




Figure 3: **Left:** The variance starts high and is reduced throughout training. From 20 epochs onwards, the variance stays stable for all classes and no overfitting can be observed. **Mid:** The bias is reduced a lot quicker than the variance, reaching its minimum at 5 epochs, and converges after 10 epochs. **Right:** The distributional correlation between training epochs shows similar to the variance that convergence happens around epoch 20. Remarkably, the ‘square’ of very high correlations indicates that the model is stable in its convergence and does not iterate through equally good solutions.




Figure 4: **Left:** Dependence between squared MMD and distributional variance. The distributional variance correlates linearly with the squared MMD. **Right:** Pearson correlation between squared MMD and distributional variance is very high ( $\approx 0.95$ ) throughout training. Approximation via deep ensembles does not deteriorate this relation. Consequently, distributional variance and kernel entropy represent viable measure of uncertainty.

used RBF kernel  $k_{\text{rbf}}(x, y) = \exp(-\gamma \|x - y\|_2^2)$ , where  $x$  and  $y$  are flattened images and  $\gamma$  a normalization factor based on the number of pixels (Schölkopf, 1997; Schölkopf & Smola, 2002; Han et al., 2012; Gretton et al., 2012a; Li et al., 2015; Bińkowski et al., 2018; Liu et al., 2020). We first analyse the generalization behavior and then assess different approaches for uncertainty estimation.

In Figure 3, we plot the distributional variance, bias, and distributional correlation throughout training for every second epoch. As can be seen, the model converges quicker for the bias than the variance (around epochs 10 and 20). Further, no overfitting occurs since both variance and bias stay small. The correlation matrix also shows convergence around epoch 20, but, more interestingly, it shows a square of high correlations in the lower right corner. This is an indication that the diffusion model training is stable in its convergence and does not iterate through different minima in the optimization landscape.

In Figure 4, we compare the relations between kernel score and  $\text{MMD}^2$  to distributional variance and predictive kernel entropy for each class. As can be seen, the predictive kernel entropy correlates strongly linearly with the generalization error (kernel score) but not so much with the generalization discrepancy ( $\text{MMD}^2$ ), while the distributional variance correlates strongly linearly with the  $\text{MMD}^2$  but not the kernel score. This corre-




Figure 5:  $\text{MMD}^2$ , variance, and bias for class ‘0’ throughout training with reduced training set of ‘0’s. After 5 epochs, mode collapse occurs, which is only expressed in the increased bias. This indicates, that mode collapse is a contrary phenomenon to overfitting.




Figure 6: **Left:** Kernel score, predictive kernel entropy, and their Pearson correlation throughout training for Glow-TTS. The entropy indicates that the model initially predicts too narrow distributions, which widen until convergence at around 10000 iterations. After convergence, the correlation between predictive entropy and kernel score is very high. **Mid:** The variance of a Deep Ensemble is also initially very small and converges at the same time as kernel score and entropy. **Right:** Comparing kernel score and kernel entropy as well as variance for 100 test instances at 16000 training iterations. Again, the correlation shows strong linearity.

lation can be observed throughout the whole training and gives a very high Pearson correlation coefficient of around 0.95. Importantly for practical settings, the correlation between  $\text{MMD}^2$  and distributional variance does not deteriorate when we use a deep ensemble trained on a single dataset.

In summary, these results demonstrate that both distributional variance as well as predictive kernel entropy are viable measure of uncertainty to predict the correctness of generated instances, either in terms of kernel score or  $\text{MMD}^2$ .

We next set out to use our estimators for bias and variance to elucidate the phenomenon of mode collapse. When generative models are used to learn the distribution of a given training set, there are often groups of different sizes present. In these cases, a common occurrence is mode collapse towards the majority groups, i.e. the model catastrophically fails to model the minority groups. To simulate this scenario in our setting, we repeat our evaluation but reduce the frequency of images of digit '0' to  $\approx 1\%$  in each training set. The bias-variance curves of digit '0' across training (Figure 5) reveals the expected mode collapse, and, most importantly, demonstrates that it is only expressed in terms of the bias. The variance term is further reduced throughout training as if no collapse occurred. This suggests that mode collapse may be seen as a contrary phenomenon to overfitting. While in overfitting, the variance increases and the bias reduces, this is vice versa for mode collapse for prolonged training.

## 5.2 AUDIO GENERATION

We next evaluate the fitting and generalization behavior of the generative flow model Glow-TTS (Kim et al., 2020) on the text-to-speech dataset LJSpeech (Ito & Johnson, 2017) throughout training (Appendix C). We train a deep ensemble via  $n = 10$  different weight initializations on 90% of the available data. We evaluate the models every 2000 training iterations by generating  $m = 10$  speech waveforms for each of 100 test instances. The evaluation includes the kernel score, kernel entropy and the distributional variance. Note that ensemble members are only required to compute the distributional variance – we compute kernel entropy for a single model, as before. Here, we use the Laplacian kernel  $k_{\text{lap}}(x, y) = \exp(-\gamma \|x - y\|_1)$  (Schölkopf & Smola, 2002), where  $x$  and  $y$  are waveforms represented by vectors and  $\gamma$  a normalization constant based on the waveform length. The results are depicted in Figure 6. Our analyses reveal that similar as for image generation, overfitting does not occur for prolonged training and the Pearson correlation between kernel entropy and kernel score is very high after convergence. But, the entropy is initially very small and increases until convergence. This indicates that a successful training requires to widen the predicted distribution. Consequently, the correlation between kernel entropy and kernel score is negative, since badly fitted instances have a more narrow predicted distribution. We also conducted the evaluations with the RBF kernel, which gives similar but slightly more erratic curves than the Laplacian kernel (Appendix C).






Figure 7: Area-Under-Curve of answer accuracy based on thresholds for different uncertainty measures. The kernel entropy outperforms other baselines in predicting the correctness of generated answers across a wide range of differently sized models.

### 5.3 NATURAL LANGUAGE GENERATION

An instance-level uncertainty measure is supposed to predict the correctness of an individual prediction and should therefore be highly correlated to the loss. In all experiments so far, we observed a very high correlation between kernel entropy and kernel score. This indicates that kernel entropy is an excellent measure of uncertainty. In the following, we examine kernel entropy to predict the correctness of LLMs on question answering datasets. Here, the setup differs from the previous experiments by two aspects.

First, we follow Kuhn et al. (2023) and do not use a kernel score but a binarized version of the RougeL similarity as loss. For two sequences  $s, t$ , it is defined as  $\text{RougeL}(s, t) = \frac{2}{\text{length}(s) + \text{length}(t)} \text{LCS}(s, t)$ , where LCS is the length of the longest common sequence between its two inputs (Lin & Och, 2004). Kuhn et al. (2023) propose to use the binary loss  $L(\text{answer}, \text{target}) = \mathbb{1}\{\text{RougeL}(\text{answer}, \text{target}) > 0.3\}$ . This turns predicting the loss value into a binary classification problem. Consequently, the AUROC is more meaningful than Pearson correlation for evaluating the performance of uncertainty measures (Kadavath et al., 2022).

Second, we do not directly use the generated answers as inputs for a kernel but, instead, their vector embeddings. A well-trained vector embedder maps text into a semantically meaningful vector space in which a kernel then measures similarities (Camacho-Collados & Pilehvar, 2018).

The investigated uncertainty baselines are predictive entropy, normalised (predictive) entropy,  $p(\text{True})$ , semantic entropy, and lexical similarity (Kuhn et al., 2023). We consider uncertainty estimation for question answering predictions of the datasets CoQA (Reddy et al., 2019) with 7983 test instances and TriviaQA (Joshi et al., 2017) with 5383 test instances. We use OPT models (Zhang et al., 2022) of all available sizes except the 30 billion parameter version, which is computationally prohibitive. For our kernel entropy, we use the RBF kernel and text embeddings computed via a pretrained e5-small-v2 (Wang et al., 2022). The results are depicted in Figure 7. As can be seen, kernel entropy is the most robust approach and outperforms other baselines for uncertainty estimation in almost all cases. We can achieve further improvements in our approach when we use alternative embedders (c.f. Appendix C). The cosine similarity, which is used in natural language processing (Steinbach et al., 2000), and other kernels show similar results as the RBF kernel in Appendix C.

## 6 CONCLUSION

In this work we introduced the first bias-variance-covariance decomposition beyond the mean squared error for kernel scores. We proposed estimators for the variance and covariance terms which only require samples of the predictive distributions but not the distributions themselves. This allows to evaluate all terms in the composition for arbitrary generative models and even in the closed-source setting. We studied empirically the fitting behavior of common models for image and audio generations, and demonstrated that kernel entropy and variance are viable measures of uncertainty. Finally, we showed that kernel entropy outperforms other baselines for predicting the correctness of LLMs in question answering tasks.

## REFERENCES

- Mikołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD GANs. In *International Conference on Learning Representations*, 2018.
- Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*, volume 4. Springer, 2006.
- Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- Gavin Brown. *Diversity in neural network ensembles*. PhD thesis, Citeseer, 2004.
- Jose Camacho-Collados and Mohammad Taher Pilehvar. From word to sense embeddings: A survey on vector representations of meaning. *Journal of Artificial Intelligence Research*, 63:743–788, 2018.
- Kacper Chwiałkowski, Heiko Strathmann, and Arthur Gretton. A kernel test of goodness of fit. In *International conference on machine learning*, pp. 2606–2615. PMLR, 2016.
- A Philip Dawid. The geometry of proper scoring rules. *Annals of the Institute of Statistical Mathematics*, 59(1):77–93, 2007.
- Morris Eaton. A method for evaluating improper prior distributions. Technical report, University of Minnesota, 1981.
- Morris L Eaton, Alessandra Giovagnoli, and Paola Sebastiani. A predictive approach to the bayesian design problem with application to normal regression models. *Biometrika*, 83(1):111–125, 1996.
- Gölge Eren and The Coqui TTS Team. Coqui TTS, January 2021. URL <https://github.com/coqui-ai/TTS>.
- Marina Fomicheva, Shuo Sun, Lisa Yankovskaya, Frédéric Blain, Francisco Guzmán, Mark Fishel, Nikolaos Aletras, Vishrav Chaudhary, and Lucia Specia. Unsupervised quality estimation for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8: 539–555, 2020.
- Jerome H Friedman. Stochastic gradient boosting. *Computational statistics & data analysis*, 38(4): 367–378, 2002.
- Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007. doi: 10.1198/016214506000001437. URL <https://doi.org/10.1198/016214506000001437>.
- Arthur Gretton, Ralf Herbrich, and Alexander J Smola. The kernel mutual information. In *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings.(ICASSP'03).*, volume 4, pp. IV–880. IEEE, 2003.
- Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(25):723–773, 2012a. URL <http://jmlr.org/papers/v13/gretton12a.html>.
- Arthur Gretton, Dino Sejdinovic, Heiko Strathmann, Sivaraman Balakrishnan, Massimiliano Pontil, Kenji Fukumizu, and Bharath K Sriperumbudur. Optimal kernel choice for large-scale two-sample tests. *Advances in neural information processing systems*, 25, 2012b.
- Sebastian Gruber and Florian Buettner. Uncertainty estimates of predictions via a general bias-variance decomposition. In *International Conference on Artificial Intelligence and Statistics*, pp. 11331–11354. PMLR, 2023.
- Shunjie Han, Cao Qubo, and Han Meng. Parameter selection in svm with rbf kernel function. In *World Automation Congress 2012*, pp. 1–4. IEEE, 2012.
- Trevor Hastie, Robert Tibshirani, Jerome H Friedman, and Jerome H Friedman. *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer, 2009.

- Achim Hekler, Titus J Brinker, and Florian Buettner. Test time augmentation meets post-hoc calibration: uncertainty quantification under real-world conditions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 14856–14864, 2023.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Keith Ito and Linda Johnson. The lj speech dataset. <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1601–1611, 2017.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*, 2022.
- Enkelejda Kasneci, Kathrin Seßler, Stefan Küchemann, Maria Bannert, Daryna Dementieva, Frank Fischer, Urs Gasser, Georg Groh, Stephan Günemann, Eyke Hüllermeier, et al. Chatgpt for good? on opportunities and challenges of large language models for education. *Learning and individual differences*, 103:102274, 2023.
- Jaehyeon Kim, Sungwon Kim, Jungil Kong, and Sungroh Yoon. Glow-tts: A generative flow for text-to-speech via monotonic alignment search. *Advances in Neural Information Processing Systems*, 33:8067–8077, 2020.
- Jonas Kübler, Wittawat Jitkrittum, Bernhard Schölkopf, and Krikamol Muandet. Learning kernel tests without data splitting. *Advances in Neural Information Processing Systems*, 33:6245–6255, 2020.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. In *The Eleventh International Conference on Learning Representations*, 2023.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
- Yujia Li, Kevin Swersky, and Rich Zemel. Generative moment matching networks. In *International conference on machine learning*, pp. 1718–1727. PMLR, 2015.
- Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. Towards general text embeddings with multi-stage contrastive learning. *arXiv preprint arXiv:2308.03281*, 2023.
- Chin-Yew Lin and Franz Josef Och. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL-04)*, pp. 605–612, 2004.
- Feng Liu, Wenkai Xu, Jie Lu, Guangquan Zhang, Arthur Gretton, and Danica J Sutherland. Learning deep kernels for non-parametric two-sample tests. In *International conference on machine learning*, pp. 6316–6326. PMLR, 2020.
- Yong Liu and Xin Yao. Ensemble learning via negative correlation. *Neural networks*, 12(10): 1399–1404, 1999a.
- Yong Liu and Xin Yao. Simultaneous training of negatively correlated neural networks in an ensemble. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 29(6):716–725, 1999b.

- Gaëlle Loosli, Stéphane Canu, and Léon Bottou. Training invariant support vector machines using selective sampling. In Léon Bottou, Olivier Chapelle, Dennis DeCoste, and Jason Weston (eds.), *Large Scale Kernel Machines*, pp. 301–320. MIT Press, Cambridge, MA., 2007. URL <http://leon.bottou.org/papers/loosli-canu-bottou-2006>.
- Andrey Malinin and Mark Gales. Uncertainty estimation in autoregressive structured prediction. In *International Conference on Learning Representations*, 2020.
- Bertalan Meskó and Eric J Topol. The imperative for regulatory oversight of large language models (or generative ai) in healthcare. *npj Digital Medicine*, 6(1):120, 2023.
- Kevin P. Murphy. *Probabilistic Machine Learning: An introduction*. MIT Press, 2022. URL [probml.ai](http://probml.ai).
- Yishuang Ning, Sheng He, Zhiyong Wu, Chunxiao Xing, and Liang-Jie Zhang. A review of deep learning based speech synthesis. *Applied Sciences*, 9(19):4050, 2019.
- OpenAI. Gpt-4 technical report, 2023.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 32*, 2019.
- Debleena Paul, Gaurav Sanap, Snehal Shenoy, Dnyaneshwar Kalyane, Kiran Kalia, and Rakesh K Tekade. Artificial intelligence in drug discovery and development. *Drug discovery today*, 26(1):80, 2021.
- Alexandre Rame, Matthieu Kirchmeyer, Thibaud Rahier, Alain Rakotomamonjy, Patrick Gallinari, and Matthieu Cord. Diverse weight averaging for out-of-distribution generalization. *Advances in Neural Information Processing Systems*, 35:10821–10836, 2022.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pp. 8821–8831. PMLR, 2021.
- Siva Reddy, Danqi Chen, and Christopher D Manning. Coqa: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266, 2019.
- Carl-Erik Särndal, Bengt Swensson, and Jan Wretman. *Model assisted survey sampling*. Springer Science & Business Media, 2003.
- Bernhard Schölkopf. *Support vector learning*. PhD thesis, Oldenbourg München, Germany, 1997.
- Bernhard Schölkopf and Alexander J Smola. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.
- Antonin Schrab, Ilmun Kim, Benjamin Guedj, and Arthur Gretton. Efficient aggregated kernel tests using incomplete  $u$ -statistics. *Advances in Neural Information Processing Systems*, 35: 18793–18807, 2022.
- Antonin Schrab, Ilmun Kim, Mélisande Albert, Béatrice Laurent, Benjamin Guedj, and Arthur Gretton. Mmd aggregated two-sample test. *Journal of Machine Learning Research (JMLR)*, 24, 2023.
- Jun Shao. *Mathematical statistics*. Springer Science & Business Media, 2003.
- Shubhanshu Shekhar, Ilmun Kim, and Aaditya Ramdas. A permutation-free kernel two-sample test. *Advances in Neural Information Processing Systems*, 35:18168–18180, 2022.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. Mpnet: Masked and permuted pre-training for language understanding. *Advances in Neural Information Processing Systems*, 33: 16857–16867, 2020.

- Michael Steinbach, George Karypis, and Vipin Kumar. A comparison of document clustering techniques. Technical report, Department of Computer Science and Engineering, University of Minnesota, 2000.
- Ingo Steinwart and Johanna F Ziegel. Strictly proper kernel scores and characteristic kernels on compact spaces. *Applied and Computational Harmonic Analysis*, 51:510–542, 2021.
- Naonori Ueda and Ryohei Nakano. Generalization error of ensemble estimators. In *Proceedings of International Conference on Neural Networks (ICNN'96)*, volume 1, pp. 90–95. IEEE, 1996.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. Text embeddings by weakly-supervised contrastive pre-training. *arXiv preprint arXiv:2212.03533*, 2022.
- Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in Neural Information Processing Systems*, 33:5776–5788, 2020.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, pp. 38–45, 2020.
- Junyi Wu and Shari Shang. Managing uncertainty in ai-enabled decision making and achieving sustainability. *Sustainability*, 12(21):8758, 2020.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.

## A OVERVIEW

In the following, we offer further details and additional experiments in Appendix C, and we provide all missing proofs in Appendix D.

## B DECOMPOSITION IN REPRODUCING KERNEL HILBERT SPACES

The literature on MMD expanded to a significant size in the last decade (Gretton et al., 2012a;b; Chwialkowski et al., 2016; Liu et al., 2020; Kübler et al., 2020; Shekhar et al., 2022; Schrab et al., 2022; 2023). The MMD is usually used in the context of reproducing kernel Hilbert spaces (RKHS). In the following, we express Theorem 3.2 according to RKHS and MMD to offer an alternative perspective on our result. Assume the kernel  $k$  is associated with an RKHS  $\mathcal{H}$  with inner product  $\langle \cdot, \cdot \rangle_{\mathcal{H}}$  and norm  $\|\cdot\|_{\mathcal{H}}$ . The norm based on  $k$  in the distribution space relates to the RKHS norm via  $\|Q\|_k = \|\mu_Q\|_{\mathcal{H}}$  with mean embedding  $\mu_Q := \mathbb{E}[k(Y, \cdot)] \in \mathcal{H}$  for a  $Y \sim Q \in \mathcal{P}$ . Consequently, given a prediction  $\hat{P}$  we have

$$\underbrace{\mathbb{E} \left[ \text{MMD}_k^2(\hat{P}, Q) \right]}_{\text{Generalization Discrepancy}} = \underbrace{\mathbb{E} \left[ \|\mu_{\hat{P}} - \mu_Q\|_{\mathcal{H}}^2 \right]}_{\text{Bias}} + \underbrace{\mathbb{E} \left[ \|\mu_{\hat{P}} - \mathbb{E}[\mu_{\hat{P}}]\|_{\mathcal{H}}^2 \right]}_{\text{Variance}}. \quad (11)$$

The covariance decomposition can be expressed similarly since  $\langle P | k | Q \rangle = \langle \mu_P, \mu_Q \rangle_{\mathcal{H}}$ . Note that the bias and variance terms in Theorem 3.2 and Equation 11 are equal.

## C EXTENDED EXPERIMENTS

In this section, we give more details on the experiments and show further results.

### C.1 EXPERIMENTAL DETAILS

We give some additional details on the experimental setup. **First, we formalize computing the kernel entropy in Algorithm 1.** Second, we describe in more detail the image, audio, and natural language generation experiments.

---

#### Algorithm 1 Uncertainty estimation via kernel entropy

---

**Require:** Generated outputs  $a_1, \dots, a_n$  for a given input, p.s.d. kernel  $k$ , (optional) embedder  $\phi$   
**if**  $\phi$  is given **then**  
  **for**  $i \in \{1, \dots, n\}$  **do**  
     $a_i \leftarrow \phi(a_i)$   
  **end for**  
**end if**  
**return**  $-\frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j \neq i}^n k(a_i, a_j)$  ▷ Estimated kernel entropy

---

#### C.1.1 IMAGE GENERATION

We use the following procedure for the simulation in Figure 2. First, we sample 32 distinct training sets from InfMNIST of size 60,000. Then, we train a conditional diffusion model on each training set for 20 epochs. We adopt implementation and hyperparameters from open source PyTorch code of a conditional diffusion model trained on normal MNIST (Paszke et al., 2019). This includes an initial learning rate of 1e-4, a batch size of 256, 400 diffusion steps, and a feature dimension of 128. We then generate 100 images of class '0' of each model after training. Finally, to get the approximate standard deviation of each tick and each line in Figure 2, we estimate the distributional variance 1000 times on randomly drawn samples without replacement of all generated images. The normalization constant in the RBF kernel is set to  $\gamma = \frac{1}{728}$  since each image has 728 pixels. We also repeated the whole procedure for other classes with similar results.




Figure 8: **Left:** Kernel score throughout training. The kernel score cannot be compared meaningfully between classes since each optimum depends on a constant specific to each target class. **Mid:**  $MMD^2$  throughout training. Here, it is easier to compare the errors since the  $MMD^2$  is zero for the optimal prediction. **Right:** The predictive kernel entropy of each class fluctuates throughout training but stays constant upon convergence.

The results in Figure 3 and 4 are produced in a similar manner. Only difference here is that we train on only 20 training sets for 40 epochs and generate only 20 images per class. We chose these numbers based on the insights gained by the previous simulation experiment. All training was done on Nvidia RTX5000 GPUs.

### C.1.2 AUDIO GENERATION

For the audio experiments, we use an implementation of Glow-TTS given in the TTS library (Eren & The Coqui TTS Team, 2021). The LJSpeech dataset consists of 13,100 instances of text-speech pairs. Each speech is of a single woman reading out loud the corresponding text. We use a random 90% of the data for training and a batch size of 32. On this single training set, we train 12 randomly initialized models for 100 epochs with an initial learning rate of  $1e-2$ . The evaluation happens every 2,000 gradient descent iterations for each model. For a single model in a single evaluation step, we generate 10 waveforms (speeches) for each of 100 test instances. The normalization constant in the Laplacian and RBF kernel is set to  $\gamma = \frac{1}{\lambda_{iter,inst}}$ , where  $\lambda_{iter,inst}$  is the longest generated waveform in each evaluation step iter and each test instance inst. **The generated audio instances are of various length, so we pad them with zeros to match their length.**

Further, all training was done on Nvidia RTX5000 GPUs. But, noteworthy to this experiment, storing the model iterations and generated waveforms required up to 400 GB of hard disk storage.

### C.1.3 NATURAL LANGUAGE GENERATION

For the natural language experiments, we adopted the experimental setup of Kuhn et al. (2023). We used their provided code implementations including the hyperparameters. This includes a temperature of  $T = 0.5$  for generating the answers used for uncertainty estimation. Similarly, we use 10 answer generations for each prompt for TriviaQA and 20 answer generations for CoQA. All natural language models are pretrained and downloaded from HuggingFace (Wolf et al., 2020). We used a single Nvidia A6000 GPU for the natural language experiments.

## C.2 ADDITIONAL RESULTS

In the following, we give some additional results for image, audio, and language generation.

### C.2.1 IMAGE GENERATION

We start with the image experiments. In Figure 8, we show the corresponding generalization error (kernel score),  $MMD^2$ , and predictive kernel entropy values throughout training of the same setup as in Figure 3. As can be seen,  $MMD^2$  is more interpretative for comparing different classes. But, the  $MMD$  can also not be evaluated in a lot of practical cases including the audio and natural language settings in this work. Further, the kernel entropy does not show a trend throughout training contrary to the audio setting seen in Figure 6.




Figure 9: Generalization curves of the digit, which we reduce in frequency during training of DDPMs on InfiMNIST. The reduced digit has only approx. 60 training instances while the other digits have approx. 6000 each. **Left:** We reduce and evaluate exclusively digit '2'. **Right:** We reduce and evaluate exclusively digit '3'. We cannot observe a mode collapse as with digit '0' in Figure 5. The bias is still elevated compared to normal digit frequencies (c.f. Figure 3).




Figure 10: Average WAV length of generated audio instances by Glow-TTS. Initially, the model generates only very short instances, which explains the low initial kernel entropy.

Further, we perform similar experiments as in Figure 5, where we observed a mode collapse. In Figure 9, we repeat the evaluation but use digit '2' and digit '3' instead of digit '0'. As can be seen via our bias-variance decomposition, mode collapse does not occur, but the bias term is still larger than in Figure 3. This underlines that our decomposition is a useful tool to investigate and analyze the fitting behaviour of generative models.

### C.2.2 AUDIO GENERATION

We repeat the evaluations of the audio generations in Section 5.2, where we used the Laplacian kernel, with the RBF kernel. It is known that the RBF kernel does not scale well to higher dimensions (Bińkowski et al., 2018). We represent the generated audio instances via vectors of various lengths, often exceeding 100,000 dimensions (c.f. Figure 10).

Consequently, we expect the RBF kernel to behave more erratic than the Laplacian kernel in the main paper. The results are shown in Figure 11.

### C.2.3 NATURAL LANGUAGE GENERATION

Next, we continue with the natural language experiments.

In Figure 12, we confirm that the correlation between the predictive kernel entropy and the RougeL (which is supposed to be maximized) has the same sign as in the image experiments in Figure 4.

We also evaluate different embedders. For this, we compare different ones, which have been pretrained on a variety of different training sets and with different embedding dimensions available on HuggingFace. These are e5-small-v2 (384 dimensions; used in the main paper) (Wang et al., 2022), gte-large (1024 dimensions) (Li et al., 2023), all-mpnet-base-v2 (768 dimensions) (Song et al., 2020), as well as all-MiniLM-l6-v2 and all-MiniLM-L12-v2 (both 384 dimensions) (Wang et al., 2020). The models all-mpnet-base-v2, all-MiniLM-l6-v2, and all-MiniLM-L12-v2 included the training set of TriviaQA in their training data, while e5-small-v2 and gte-large did not. Consequently, these three models are an unfair comparison for TriviaQA to other baselines not using its training set. In Figure 13, we compare the ability of each embedder in combination with the RBF kernel to predict






Figure 11: Audio generation results for kernel score, kernel entropy, and negative Pearson correlation between them with the RBF kernel. The trends are similar as in Figure 6 but more erratic and the absolute correlation is slightly less.




Figure 12: Kernel density estimation of kernel entropy and RougeL between answer and target for opt-1.3b model on CoQA. A large RougeL corresponds to a high accuracy and a low error. Consequently, low kernel entropy indicates a high likelihood of answer correctness.

the answer accuracy for the same settings as in Figure 7 in the main paper. As we can see for CoQA, all embedders provide approximately similar performance. This indicates that kernel entropy is a robust approach as long as the embedder is meaningful. The results for TriviaQA are similar, but all-MiniLM-L6-v2 and all-MiniLM-L12-v2 perform comparably better. We hypothesis that this is due to them being trained on the training set of TriviaQA. This suggests that we can achieve even better uncertainty estimates by using a task specific training set.

We also compare the impact of different kernel choices in Figure 14. The differences are marginal for cosine similarity, and RBF and polynomial kernel. Contrary, the Laplacian kernel performs often worse. The lack of a difference between RBF or polynomial kernel and cosine similarity is surprising considering that e5-small-v2 has been trained using the cosine similarity (Wang et al., 2022). Our results suggests that the embedder has substantially more influence than the kernel, and that resources should be spent on optimizing the former and not the latter.

Last, we compare different number of generated answers to estimate the kernel entropy in the case of Opt-13b on CoQA in Figure 15. As can be seen, more samples are continuously better. Consequently, we expect to get even better results in Figure 7 for larger numbers of generations.

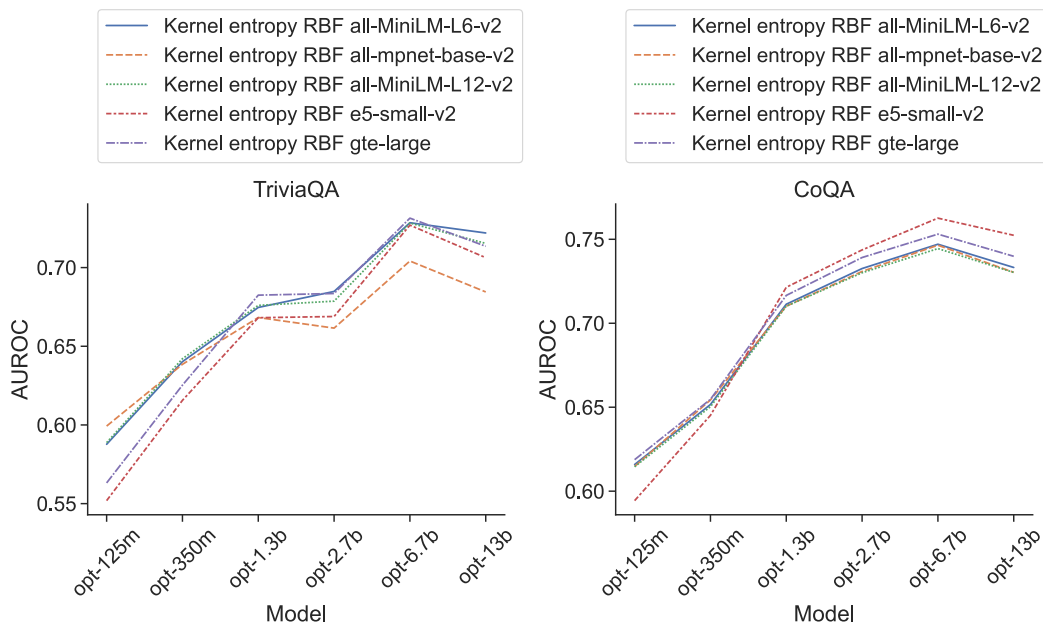


Figure 13: Comparing uncertainty estimates of kernel entropy for different embedders and RBF kernel. In both cases, the differences are not substantial. But, all-MiniLM-L6-v2 and all-MiniLM-L12-v2 perform comparably better on TriviaQA, likely due to them being trained on the TriviaQA training set.

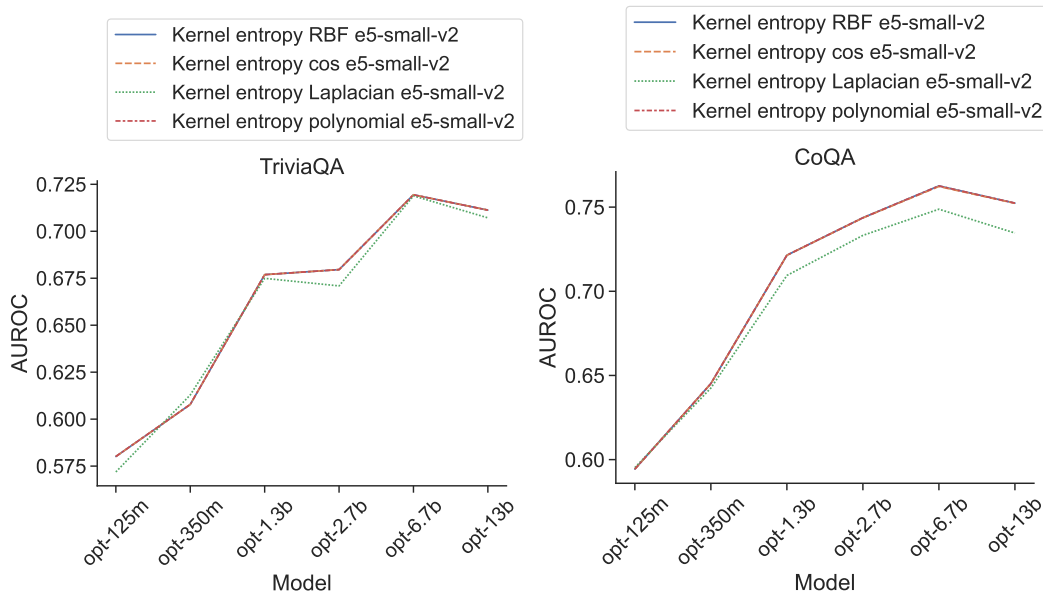


Figure 14: Comparing uncertainty estimates of kernel entropy for cosine similarity and RBF, polynomial, and Laplacian kernel. Even though the embedder e5-small-v2 is trained via cosine similarity, the performance differences are marginal.




Figure 15: Number of generated answers to estimate kernel entropy compared to the AUROC for Opt-13b on CoQA. More generations monotonically improve kernel entropy as an uncertainty estimate.

## D MISSING PROOFS

In this section, we give all missing proofs for Theorem 3.2 and for the statements in Section 4. We first start with the proof for the bias-variance-covariance decomposition in Section D.1. Then in Section D.2, we solve the expectation and the variance of the distributional variance estimator proposed in Equation 7. Last, we do the same for the distributional covariance estimator of Equation 8 in Section D.3.

### D.1 BIAS-VARIANCE-COVARIANCE DECOMPOSITION

In Theorem 3.2, the covariance decomposition is introduced after stating the more simpler bias-variance decomposition. Here, we prove both in one go. Assume we have target  $Y \sim Q \in \mathcal{P}$  and ensemble prediction  $\hat{P}^{(n)} := \frac{1}{n} \sum_{i=1}^n \hat{P}_i$  of  $n \in \mathbb{N}$  identically distributed predictions  $\hat{P}_1, \dots, \hat{P}_n$  with outcomes in  $\mathcal{P}$ . Further, assume that all expectations in the following are finite. We will use  $\hat{P} := \hat{P}_1$  and  $\hat{P}' := \hat{P}_2$ . Then, the decomposition can be constructed via

$$\begin{aligned}
& \mathbb{E} \left[ -S_k \left( \hat{P}^{(n)}, Y \right) \right] \\
&= \mathbb{E} \left[ \left\| \hat{P}^{(n)} \right\|_k^2 - 2 \langle \hat{P}^{(n)} \mid k \mid \delta_Y \rangle \right] \\
&\stackrel{(i)}{=} \mathbb{E} \left[ \left\| \hat{P}^{(n)} \right\|_k^2 - 2 \langle \hat{P}^{(n)} \mid k \mid Q \rangle \right] \\
&= \mathbb{E} \left[ \left\| \hat{P}^{(n)} \right\|_k^2 - 2 \langle \mathbb{E} [\hat{P}^{(n)}] \mid k \mid Q \rangle \right] \\
&= \mathbb{E} \left[ \left\| \hat{P}^{(n)} \right\|_k^2 - 2 \langle \mathbb{E} [\hat{P}^{(n)}] \mid k \mid Q \rangle + 2 \left\| \mathbb{E} [\hat{P}^{(n)}] \right\|_k^2 - 2 \mathbb{E} \left[ \langle \mathbb{E} [\hat{P}^{(n)}] \mid k \mid \hat{P}^{(n)} \rangle \right] \right] \\
&= \mathbb{E} \left[ \left\| \hat{P}^{(n)} - \mathbb{E} [\hat{P}^{(n)}] \right\|_k^2 - 2 \langle \mathbb{E} [\hat{P}^{(n)}] \mid k \mid Q \rangle + \left\| \mathbb{E} [\hat{P}^{(n)}] \right\|_k^2 \right] \\
&= \mathbb{E} \left[ \left\| \hat{P}^{(n)} - \mathbb{E} [\hat{P}^{(n)}] \right\|_k^2 - 2 \langle \mathbb{E} [\hat{P}^{(n)}] \mid k \mid Q \rangle + \left\| \mathbb{E} [\hat{P}^{(n)}] \right\|_k^2 + \|Q\|_k^2 - \|Q\|_k^2 \right] \\
&= -\|Q\|_k^2 + \left\| \mathbb{E} [\hat{P}^{(n)}] - Q \right\|_k^2 + \mathbb{E} \left[ \left\| \hat{P}^{(n)} - \mathbb{E} [\hat{P}^{(n)}] \right\|_k^2 \right] \\
&\stackrel{(ii)}{=} -\|Q\|_k^2 + \left\| \mathbb{E} [\hat{P}] - Q \right\|_k^2 + \mathbb{E} \left[ \left\| \hat{P}^{(n)} - \mathbb{E} [\hat{P}] \right\|_k^2 \right] \\
&= -\|Q\|_k^2 + \left\| \mathbb{E} [\hat{P}] - Q \right\|_k^2 + \mathbb{E} \left[ \left\| \frac{1}{n} \sum_{i=1}^n \hat{P}_i - \mathbb{E} [\hat{P}] \right\|_k^2 \right] \\
&= -\|Q\|_k^2 + \left\| \mathbb{E} [\hat{P}] - Q \right\|_k^2 + \frac{1}{n^2} \sum_{i,j=1}^n \mathbb{E} \left[ \langle \hat{P}_i - \mathbb{E} [\hat{P}] \mid k \mid \hat{P}_j - \mathbb{E} [\hat{P}] \rangle \right] \\
&= -\|Q\|_k^2 + \left\| \mathbb{E} [\hat{P}] - Q \right\|_k^2 + \frac{1}{n^2} \sum_{i,j=1}^n \text{Cov}_k (\hat{P}_i, \hat{P}_j) \\
&= -\|Q\|_k^2 + \frac{1}{n^2} \sum_{i=1}^n \text{Var}_k (\hat{P}_i) + \frac{1}{n^2} \sum_{\substack{i,j=1 \\ i \neq j}}^n \text{Cov}_k (\hat{P}_i, \hat{P}_j) + \left\| \mathbb{E} [\hat{P}] - Q \right\|_k^2 \\
&\stackrel{(ii)}{=} -\|Q\|_k^2 + \left\| \mathbb{E} [\hat{P}] - Q \right\|_k^2 + \frac{1}{n^2} n \text{Var}_k (\hat{P}) + \frac{1}{n^2} n(n-1) \text{Cov}_k (\hat{P}, \hat{P}') \\
&= \underbrace{-\|Q\|_k^2}_{\text{Noise}} + \underbrace{\left\| \mathbb{E} [\hat{P}] - Q \right\|_k^2}_{\text{Bias}} + \underbrace{\frac{1}{n} \text{Var}_k (\hat{P})}_{\text{Variance}} + \underbrace{\frac{n-1}{n} \text{Cov}_k (\hat{P}, \hat{P}')}_{\text{Covariance}}
\end{aligned} \tag{12}$$

(i)  $Y$  and  $\hat{P}^{(n)}$  independently distributed

(ii)  $\hat{P}_1, \dots, \hat{P}_n$  identically distributed

## D.2 DISTRIBUTIONAL VARIANCE ESTIMATOR

In general, the notation of Eaton (1981) allows us to write for random variables  $P$  and  $Q$  with outcomes in a distribution space, and for independent  $X \sim P$  and  $Y \sim Q$  that

$$\mathbb{E} [k(X, Y) \mid P, Q] = \int \int k(x, y) dP \otimes Q(x, y) = \int \int k(x, y) dP(x) dQ(y) = \langle P \mid k \mid Q \rangle, \tag{13}$$

where we used Tonelli's Theorem to split the integral.

Note we have for  $j \neq t$  that  $X_{ij}$  and  $X_{it}$  are independent given  $P_i$  for all  $i = 1, \dots, n$ . From this follows

$$\begin{aligned}
\mathbb{E}[k(X_{ij}, X_{it})] &= \mathbb{E}[\mathbb{E}[k(X_{ij}, X_{it}) | P_i]] \\
&\stackrel{\text{iid}}{=} \mathbb{E}[\mathbb{E}[k(X_{i1}, X_{i2}) | P_i]] \\
&\stackrel{\text{iid}}{=} \mathbb{E}[\langle P_i | k | P_i \rangle] \\
&\stackrel{\text{iid}}{=} \mathbb{E}[\langle P | k | P \rangle]
\end{aligned} \tag{14}$$

as well as for  $i \neq s$

$$\begin{aligned}
\mathbb{E}[k(X_{ij}, X_{st})] &= \mathbb{E}[\mathbb{E}[k(X_{ij}, X_{st}) | P_i, P_s]] \\
&\stackrel{\text{iid}}{=} \mathbb{E}[\mathbb{E}[k(X_{i1}, X_{s1}) | P_i, P_s]] \\
&\stackrel{\text{iid}}{=} \mathbb{E}[\langle P_i | k | P_s \rangle] \\
&\stackrel{\text{iid}}{=} \langle \mathbb{E}[P] | k | \mathbb{E}[P] \rangle.
\end{aligned} \tag{15}$$

### D.2.1 EXPECTATION OF THE ESTIMATOR

Using these equations gives

$$\begin{aligned}
&\mathbb{E} \left[ \widehat{\text{Var}}_k^{(n,m)} \right] \\
&= \mathbb{E} \left[ \frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m \left( \frac{1}{m-1} \sum_{\substack{t=1 \\ t \neq j}}^m k(X_{ij}, X_{it}) - \frac{1}{(n-1)m} \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{t=1}^m k(X_{ij}, X_{st}) \right) \right] \\
&= \frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m \left( \frac{1}{m-1} \sum_{\substack{t=1 \\ t \neq j}}^m \mathbb{E}[k(X_{ij}, X_{it})] - \frac{1}{(n-1)m} \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{t=1}^m \mathbb{E}[k(X_{ij}, X_{st})] \right) \\
&\stackrel{\text{Eq. 14 \& 15}}{=} \frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m \left( \frac{1}{m-1} \sum_{\substack{t=1 \\ t \neq j}}^m \mathbb{E}[\langle P | k | P \rangle] - \frac{1}{(n-1)m} \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{t=1}^m \langle \mathbb{E}[P] | k | \mathbb{E}[P] \rangle \right) \\
&= \mathbb{E}[\langle P | k | P \rangle] - \langle \mathbb{E}[P] | k | \mathbb{E}[P] \rangle \\
&= \mathbb{E}[\langle P - \mathbb{E}[P] | k | P - \mathbb{E}[P] \rangle] \\
&= \text{Var}_k[P]
\end{aligned} \tag{16}$$

### D.2.2 VARIANCE OF THE ESTIMATOR

Note that for a U-statistic  $\hat{U}_n = \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{\substack{j=1 \\ j \neq i}}^n h(X_i, X_j)$  based on i.i.d. samples  $X_1, \dots, X_n$  and symmetric kernel  $h$ , the estimator variance is given by

$$\mathbb{V}(\hat{U}_n) = \frac{2}{n(n-1)} \mathbb{V}(h(X_1, X_2)) + \frac{4(n-2)}{n(n-1)} \mathbb{V}(\mathbb{E}[h(X_1, X_2) | X_2]) \tag{17}$$

(Shao, 2003).

We will use the law of total variance (TV) several times to create independence between the summands, since for dependent random variables  $X$  and  $Y$  with scalars  $a, b$  we have  $\mathbb{V}(aX + bY) = a^2\mathbb{V}(X) + b^2\mathbb{V}(Y) + 2ab \text{Cov}(X, Y)$  (Shao, 2003).

From this also follows that

$$\begin{aligned}
& \mathbb{V} \left( \widehat{\text{Var}}_k^{(n,m)} \right) \\
&= \mathbb{V} \left( \frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m \left( \frac{1}{m-1} \sum_{\substack{t=1 \\ t \neq j}}^m k(X_{ij}, X_{it}) - \frac{1}{(n-1)m} \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{t=1}^m k(X_{ij}, X_{st}) \right) \right) \\
&= \mathbb{V} \left( \frac{1}{nm(m-1)} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m k(X_{ij}, X_{it}) \right) + \mathbb{V} \left( \frac{1}{n(n-1)m^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{t=1}^m k(X_{ij}, X_{st}) \right) \\
&\quad - 2 \text{Cov} \left( \frac{1}{nm(m-1)} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m k(X_{ij}, X_{it}), \frac{1}{n(n-1)m^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{t=1}^m k(X_{ij}, X_{st}) \right)
\end{aligned} \tag{18}$$

We will analyse each term successively and then combine the results further down in equation 27.

$$\begin{aligned}
& \mathbb{V} \left( \frac{1}{nm(m-1)} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m k(X_{ij}, X_{it}) \right) \\
&\stackrel{\text{TV}}{=} \mathbb{V} \left( \mathbb{E} \left[ \frac{1}{nm(m-1)} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m k(X_{ij}, X_{it}) \mid P_1 \dots P_n \right] \right) \\
&\quad + \mathbb{E} \left[ \mathbb{V} \left( \frac{1}{nm(m-1)} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m k(X_{ij}, X_{it}) \mid P_1 \dots P_n \right) \right] \\
&\stackrel{\text{iid}}{=} \mathbb{V} \left( \frac{1}{nm(m-1)} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \mathbb{E} [k(X_{ij}, X_{it}) \mid P_i] \right) \\
&\quad + \mathbb{E} \left[ \frac{1}{n^2} \sum_{i=1}^n \mathbb{V} \left( \frac{1}{m(m-1)} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m k(X_{ij}, X_{it}) \mid P_i \right) \right] \\
&\stackrel{\text{iid}}{=} \mathbb{V} \left( \frac{1}{n} \sum_{i=1}^n \mathbb{E} [k(X_{i1}, X_{i2}) \mid P_i] \right) + \frac{1}{n} \mathbb{E} \left[ \mathbb{V} \left( \frac{1}{m(m-1)} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m k(X_{1j}, X_{1t}) \mid P_1 \right) \right] \\
&\stackrel{\text{iid}}{=} \frac{1}{n} \mathbb{V} (\langle P \mid k \mid P \rangle) + \frac{1}{n} \mathbb{E} \left[ \mathbb{V} \left( \frac{1}{m(m-1)} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m k(X_{1j}, X_{1t}) \mid P_1 \right) \right] \\
&\stackrel{\text{(i)}}{=} \frac{1}{n} \mathbb{V} (\langle P \mid k \mid P \rangle) + \frac{4(m-2)}{nm(m-1)} \zeta_1 + \frac{2}{nm(m-1)} \zeta_2
\end{aligned} \tag{19}$$

(i) with  $\zeta_1 = \mathbb{E}[\mathbb{V}(\mathbb{E}[k(X_{11}, X_{12}) | X_{12}, P_1] | P_1)]$  and  $\zeta_2 = \mathbb{E}[\mathbb{V}(k(X_{11}, X_{12}) | P_1)]$  based on equation 17.

$$\begin{aligned}
& \mathbb{V} \left( \frac{1}{n(n-1)m^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{t=1}^m k(X_{ij}, X_{st}) \right) \\
& \stackrel{\text{TV}}{=} \underbrace{\mathbb{V} \left( \mathbb{E} \left[ \frac{1}{n(n-1)m^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{t=1}^m k(X_{ij}, X_{st}) \mid P_1 \dots P_n \right] \right)}_{\text{(I):=}} \\
& + \underbrace{\mathbb{E} \left[ \mathbb{V} \left( \frac{1}{n(n-1)m^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{t=1}^m k(X_{ij}, X_{st}) \mid P_1 \dots P_n \right) \right]}_{\text{(II):=}}.
\end{aligned} \tag{20}$$

Due to the length of the expression, we first solve (I) and then (II).

$$\begin{aligned}
\text{(I)} & \stackrel{\text{iid}}{=} \mathbb{V} \left( \frac{1}{n(n-1)m^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{t=1}^m \mathbb{E}[k(X_{ij}, X_{st}) \mid P_i, P_s] \right) \\
& \stackrel{\text{iid}}{=} \mathbb{V} \left( \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \langle P_i \mid k \mid P_s \rangle \right) \\
& \stackrel{\text{(i)}}{=} \frac{4(n-2)}{n(n-1)} \zeta_3 + \frac{2}{n(n-1)} \zeta_4
\end{aligned} \tag{21}$$

(i) with  $\zeta_3 = \mathbb{V}(\mathbb{E}[\langle P_1 \mid k \mid P_2 \rangle \mid P_1])$  and  $\zeta_4 = \mathbb{V}(\langle P_1 \mid k \mid P_2 \rangle)$  based on equation 17.

For the next term note that  $\frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, X_{st})$  and  $\frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{aj}, X_{bt})$  are independent given  $P_i, P_s, P_a, P_b$  for  $i \neq s \neq a \neq b$  from which follows

$$\text{Cov} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, X_{st}), \frac{1}{m^2} \sum_{c=1}^m \sum_{d=1}^m k(X_{ac}, X_{bd}) \right) = 0. \tag{22}$$

Consequently, we have

$$\begin{aligned}
(\text{II}) = & \mathbb{E} \left[ \frac{1}{n^2 (n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \mathbb{V} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, X_{st}) \mid P_i, P_s \right) \right] \\
& + \mathbb{E} \left[ \frac{1}{n^2 (n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{b=1 \\ b \neq i \\ b \neq s}}^n \text{Cov} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, X_{st}), \frac{1}{m^2} \sum_{c=1}^m \sum_{d=1}^m k(X_{ic}, X_{bd}) \mid P_i, P_s, P_b \right) \right] \\
& + \mathbb{E} \left[ \frac{1}{n^2 (n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{a=1 \\ a \neq s \\ a \neq i}}^n \text{Cov} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, X_{st}), \frac{1}{m^2} \sum_{c=1}^m \sum_{d=1}^m k(X_{ac}, X_{sd}) \mid P_i, P_s, P_a \right) \right] \\
\stackrel{\text{sym}}{=} & \underbrace{\mathbb{E} \left[ \frac{1}{n^2 (n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \mathbb{V} \left( \frac{1}{m^2} \sum_{c=1}^m \sum_{d=1}^m k(X_{ic}, X_{sd}) \mid P_i, P_s \right) \right]}_{(\text{IIa}):=} \\
& + \underbrace{\mathbb{E} \left[ \frac{2}{n^2 (n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{b=1 \\ b \neq i \\ b \neq s}}^n \text{Cov} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, X_{st}), \frac{1}{m^2} \sum_{c=1}^m \sum_{d=1}^m k(X_{ic}, X_{bd}) \mid P_i, P_s, P_b \right) \right]}_{(\text{IIb}):=}.
\end{aligned} \tag{23}$$

Due to the length of the expressions, we again first look at (IIa) and then (IIb).



$$\begin{aligned}
\text{(IIa)} &\stackrel{\text{iid}}{=} \mathbb{E} \left[ \frac{1}{n^2 (n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \mathbb{V} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, X_{st}) \mid P_i, P_s \right) \right] \\
&\stackrel{\text{iid}}{=} \frac{1}{n(n-1)} \mathbb{E} \left[ \mathbb{V} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{1j}, X_{2t}) \mid P_1, P_2 \right) \right] \\
&= \frac{1}{n(n-1)} \mathbb{E} \left[ \text{Cov} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{1j}, X_{2t}), \frac{1}{m^2} \sum_{i=1}^m \sum_{s=1}^m k(X_{1i}, X_{2s}) \mid P_1, P_2 \right) \right] \\
&= \frac{1}{n(n-1)m^4} \sum_{j=1}^m \sum_{t=1}^m \sum_{i=1}^m \sum_{s=1}^m \mathbb{E} [\text{Cov}(k(X_{1j}, X_{2t}), k(X_{1i}, X_{2s}) \mid P_1, P_2)] \\
&\stackrel{\text{iid}}{=} \frac{1}{n(n-1)m^4} \sum_{j=1}^m \sum_{t=1}^m \mathbb{E} [\text{Cov}(k(X_{1j}, X_{2t}), k(X_{1j}, X_{2t}) \mid P_1, P_2)] \\
&\quad + \frac{1}{n(n-1)m^4} \sum_{j=1}^m \sum_{t=1}^m \sum_{\substack{i=1 \\ i \neq j}}^m \mathbb{E} [\text{Cov}(k(X_{1j}, X_{2t}), k(X_{1i}, X_{2t}) \mid P_1, P_2)] \\
&\quad + \frac{1}{n(n-1)m^4} \sum_{j=1}^m \sum_{t=1}^m \sum_{\substack{s=1 \\ s \neq t}}^m \mathbb{E} [\text{Cov}(k(X_{1j}, X_{2t}), k(X_{1j}, X_{2s}) \mid P_1, P_2)] \\
&\stackrel{\text{iid}}{=} \frac{1}{n(n-1)m^2} \mathbb{E} [\text{Cov}(k(X_{11}, X_{21}), k(X_{11}, X_{21}) \mid P_1, P_2)] \\
&\quad + \frac{m-1}{n(n-1)m^2} \mathbb{E} [\text{Cov}(k(X_{11}, X_{21}), k(X_{12}, X_{21}) \mid P_1, P_2)] \\
&\quad + \frac{m-1}{n(n-1)m^2} \mathbb{E} [\text{Cov}(k(X_{11}, X_{21}), k(X_{11}, X_{22}) \mid P_1, P_2)] \\
&\stackrel{\text{(i)}}{=} \frac{1}{n(n-1)m^2} \underbrace{\mathbb{E} [\mathbb{V}(k(X_{11}, X_{21}) \mid P_1, P_2)]}_{\zeta_6 :=} \\
&\quad + \frac{2(m-1)}{n(n-1)m^2} \underbrace{\mathbb{E} [\text{Cov}(k(X_{11}, X_{21}), k(X_{12}, X_{21}) \mid P_1, P_2)]}_{\zeta_5 :=}
\end{aligned} \tag{24}$$

(i) follows from symmetry of  $k$  and assumption of identical distributions.

Further, we have  $\text{Cov}(k(X_{ij}, X_{st}), k(X_{ic}, X_{bd}) \mid P_i, P_s, P_b) = 0$  whenever  $c \neq j$  (since  $s \neq b$  and independence assumption), giving

$$\begin{aligned}
\text{(IIb)} &= \mathbb{E} \left[ \frac{2}{n^2 (n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{b=1 \\ b \neq i \\ b \neq s}}^n \text{Cov} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, X_{st}), \frac{1}{m^2} \sum_{c=1}^m \sum_{d=1}^m k(X_{ic}, X_{bd}) \mid P_i, P_s, P_b \right) \right] \\
&= \mathbb{E} \left[ \frac{2}{n^2 (n-1)^2 m^4} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{b=1 \\ b \neq i \\ b \neq s}}^n \sum_{j=1}^m \sum_{t=1}^m \sum_{d=1}^m \text{Cov} (k(X_{ij}, X_{st}), k(X_{ij}, X_{bd}) \mid P_i, P_s, P_b) \right] \\
&\stackrel{\text{iid}}{=} \mathbb{E} \left[ \frac{2}{n^2 (n-1)^2 m^4} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{b=1 \\ b \neq i \\ b \neq s}}^n \sum_{j=1}^m \sum_{t=1}^m \sum_{d=1}^m \text{Cov} (k(X_{11}, X_{21}), k(X_{11}, X_{31}) \mid P_1, P_2, P_3) \right] \\
&= \frac{2(n-2)}{n(n-1)m} \underbrace{\mathbb{E} [\text{Cov} (k(X_{11}, X_{21}), k(X_{11}, X_{31}) \mid P_1, P_2, P_3)]}_{\zeta_9 :=}
\end{aligned}$$

(25)

The only term left is

$$\begin{aligned}
& \text{Cov} \left( \frac{1}{nm(m-1)} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m k(X_{ij}, X_{it}), \frac{1}{n(n-1)m^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{t=1}^m k(X_{ij}, X_{st}) \right) \\
&= \frac{1}{n^2(n-1)m^3(m-1)} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \sum_{o=1}^n \sum_{p=1}^m \sum_{\substack{s=1 \\ s \neq o}}^n \sum_{r=1}^m \text{Cov}(k(X_{ij}, X_{it}), k(X_{op}, X_{sr})) \\
&\stackrel{\text{(i)}}{=} \frac{1}{n^2(n-1)m^3(m-1)} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \sum_{p=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{r=1}^m \\
&\quad (\text{Cov}(k(X_{ij}, X_{it}), k(X_{ip}, X_{sr})) + \text{Cov}(k(X_{ij}, X_{it}), k(X_{sp}, X_{ir}))) \\
&\stackrel{\text{(ii)}}{=} \frac{2}{n^2(n-1)m^3(m-1)} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \sum_{p=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{r=1}^m \text{Cov}(k(X_{ij}, X_{it}), k(X_{ip}, X_{sr})) \\
&\stackrel{\text{iid}}{=} \frac{2mn(n-1)}{n^2(n-1)m^3(m-1)} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \sum_{p=1}^m \text{Cov}(k(X_{1j}, X_{1t}), k(X_{1p}, X_{21})) \\
&= \frac{2}{nm^2(m-1)} \left( \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \sum_{p \in \{j, t\}} \text{Cov}(k(X_{1j}, X_{1t}), k(X_{1p}, X_{21})) + \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \sum_{\substack{p=1 \\ p \neq j \\ p \neq t}}^m \text{Cov}(k(X_{1j}, X_{1t}), k(X_{1p}, X_{21})) \right) \\
&\stackrel{\text{(iii)}}{=} \frac{2}{nm^2(m-1)} \left( \underbrace{\sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \sum_{p \in \{j, t\}} \text{Cov}(k(X_{11}, X_{12}), k(X_{11}, X_{21}))}_{2m(m-1) \text{ summands}} + \underbrace{\sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \sum_{\substack{p=1 \\ p \neq j \\ p \neq t}}^m \text{Cov}(k(X_{11}, X_{12}), k(X_{13}, X_{21}))}_{m(m-1)(m-2) \text{ summands}} \right) \\
&= \frac{4m(m-1)}{nm^2(m-1)} \text{Cov}(k(X_{11}, X_{12}), k(X_{11}, X_{21})) + \frac{2m(m-1)(m-2)}{nm^2(m-1)} \text{Cov}(k(X_{11}, X_{12}), k(X_{13}, X_{21})) \\
&= \frac{4}{nm} \underbrace{\text{Cov}(k(X_{11}, X_{12}), k(X_{11}, X_{21}))}_{\zeta_7 :=} + \frac{2(m-2)}{nm} \underbrace{\text{Cov}(k(X_{11}, X_{12}), k(X_{13}, X_{21}))}_{\zeta_8 :=}.
\end{aligned} \tag{26}$$

(i) when  $i \neq o \neq s$ , then  $\text{Cov}(k(X_{ij}, X_{it}), k(X_{op}, X_{sr})) = 0$ .

(ii) symmetry of  $k$  and iid property result in  $\text{Cov}(k(X_{ij}, X_{it}), k(X_{ip}, X_{sr})) = \text{Cov}(k(X_{ij}, X_{it}), k(X_{sp}, X_{ir}))$  as long as  $i \neq s$

(iii) iid and symmetry of  $k$ .

It follows from inserting equation 19, equation 21, equation 25, and equation 26, into equation 18 that

$$\begin{aligned}
\mathbb{V}\left(\widehat{\text{Var}}_k^{(n,m)}\right) &= \underbrace{\frac{1}{n}\mathbb{V}(\langle P | k | P \rangle)}_{\mathcal{O}\left(\frac{1}{n}\right)} + \underbrace{\frac{4(n-2)}{n(n-1)}\zeta_3 - \frac{4(m-2)}{nm}\zeta_8}_{\mathcal{O}\left(\frac{1}{n}\right)} + \underbrace{\frac{2}{n(n-1)}\zeta_4}_{\mathcal{O}\left(\frac{1}{n^2}\right)} \\
&\quad + \underbrace{\frac{4(m-2)}{nm(m-1)}\zeta_1 - \frac{8}{nm}\zeta_7 + \frac{2(n-2)}{n(n-1)m}\zeta_9}_{\mathcal{O}\left(\frac{1}{nm}\right)} \\
&\quad + \underbrace{\frac{2}{nm(m-1)}\zeta_2}_{\mathcal{O}\left(\frac{1}{nm^2}\right)} + \underbrace{\frac{2(m-1)}{n(n-1)m^2}\zeta_5}_{\mathcal{O}\left(\frac{1}{n^2m}\right)} + \underbrace{\frac{1}{n(n-1)m^2}\zeta_6}_{\mathcal{O}\left(\frac{1}{n^2m^2}\right)}.
\end{aligned} \tag{27}$$

In summary, our estimator is in  $\mathcal{O}\left(\frac{1}{n}\left(1 + \frac{1}{m}\right)\right)$  and consequently consistent w.r.t.  $n$  but not  $m$ .

### D.2.3 ILLUSTRATION OF THE ESTIMATOR

The estimator can also be visualized as the following. For  $i, s \in \{1, \dots, n\}$  define the quadratic matrices  $\mathbf{K}_{is} = (k_{is_{jt}})_{j,t=1\dots m} \in \mathbb{R}^{m \times m}$  with entries  $k_{is_{jt}} = k(X_{ij}, X_{st})$ . Then we have the colored block matrix

$$\begin{pmatrix} \mathbf{K}_{11} & \cdots & \cdots & \mathbf{K}_{1n} \\ \vdots & \ddots & \mathbf{K}_{is} & \vdots \\ \vdots & \mathbf{K}_{si} & \ddots & \vdots \\ \mathbf{K}_{1n} & \cdots & \cdots & \mathbf{K}_{nn} \end{pmatrix} = \left( \begin{array}{cccc|cc} k_{1111} & \cdots & \cdots & k_{111m} & k_{1211} & \cdots \\ \vdots & \ddots & & \vdots & \vdots & \ddots \\ \vdots & & k_{11jt} & \vdots & \vdots & \ddots \\ k_{11m1} & \cdots & k_{11tj} & \cdots & k_{11mm} & k_{12m1} \\ \hline k_{2111} & \cdots & \cdots & k_{211m} & k_{2211} & \cdots \\ \vdots & \ddots & \ddots & \vdots & \vdots & \ddots \end{array} \right). \tag{28}$$

The proposed distributional variance estimator is then the average of all cyan entries without the red ones (blocks on diagonal without diagonal entries) minus the average of all black entries (off-diagonal blocks).

### D.3 DISTRIBUTIONAL COVARIANCE ESTIMATOR

Note that we have under the given i.i.d. assumptions that

$$\begin{aligned}
\mathbb{E}[k(X_{ij}, Y_{it})] &= \mathbb{E}[\mathbb{E}[k(X_{ij}, Y_{it}) | P_i, Q_i]] \\
&\stackrel{\text{iid}}{=} \mathbb{E}[\mathbb{E}[k(X_{i1}, Y_{i1}) | P_i, Q_i]] \\
&\stackrel{\text{iid}}{=} \mathbb{E}[\langle P_i | k | Q_i \rangle] \\
&\stackrel{\text{iid}}{=} \mathbb{E}[\langle P | k | Q \rangle]
\end{aligned} \tag{29}$$

as well as for  $i \neq s$

$$\begin{aligned}
\mathbb{E}[k(X_{ij}, Y_{st})] &= \mathbb{E}[\mathbb{E}[k(X_{ij}, Y_{st}) | P_i, Q_s]] \\
&\stackrel{\text{iid}}{=} \mathbb{E}[\mathbb{E}[k(X_{i1}, Y_{s1}) | P_i, Q_s]] \\
&\stackrel{\text{iid}}{=} \mathbb{E}[\langle P_i | k | Q_s \rangle] \\
&\stackrel{\text{iid}}{=} \langle \mathbb{E}[P] | k | \mathbb{E}[Q] \rangle.
\end{aligned} \tag{30}$$

## D.3.1 EXPECTATION OF THE ESTIMATOR

Now, we can prove that the covariance estimator is unbiased, i.e.

$$\begin{aligned}
& \mathbb{E} \left[ \widehat{\text{Cov}}_k^{(n,m)}(\mathbf{X}, \mathbf{Y}) \right] \\
&= \mathbb{E} \left[ \frac{1}{nm^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m \left( k(X_{ij}, Y_{it}) - \frac{1}{n-1} \sum_{\substack{s=1 \\ s \neq i}}^n k(X_{ij}, Y_{st}) \right) \right] \\
&= \frac{1}{nm^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m \left( \mathbb{E} [k(X_{ij}, Y_{it})] - \frac{1}{n-1} \sum_{\substack{s=1 \\ s \neq i}}^n \mathbb{E} [k(X_{ij}, Y_{st})] \right) \tag{31} \\
&\stackrel{\text{Eq. 29 \& 30}}{=} \frac{1}{nm^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m \left( \mathbb{E} [\langle P | k | Q \rangle] - \frac{1}{n-1} \sum_{\substack{s=1 \\ s \neq i}}^n \langle \mathbb{E} [P] | k | \mathbb{E} [Q] \rangle \right) \\
&= \mathbb{E} [\langle P | k | Q \rangle] - \langle \mathbb{E} [P] | k | \mathbb{E} [Q] \rangle \\
&= \mathbb{E} [\langle P - \mathbb{E} [P] | k | Q - \mathbb{E} [Q] \rangle] \\
&= \text{Cov}_k(P, Q).
\end{aligned}$$

## D.3.2 VARIANCE OF THE ESTIMATOR

Similar to the variance case, we also analyse its convergence rate:

$$\begin{aligned}
& \mathbb{V} \left( \widehat{\text{Cov}}_k^{(n,m)}(\mathbf{X}, \mathbf{Y}) \right) \\
&= \mathbb{V} \left( \frac{1}{nm^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m \left( k(X_{ij}, Y_{it}) - \frac{1}{n-1} \sum_{\substack{s=1 \\ s \neq i}}^n k(X_{ij}, Y_{st}) \right) \right) \\
&= \mathbb{V} \left( \frac{1}{nm^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, Y_{it}) \right) + \mathbb{V} \left( \frac{1}{n(n-1)m^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n k(X_{ij}, Y_{st}) \right) \\
&\quad - 2 \text{Cov} \left( \frac{1}{nm^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, Y_{it}), \frac{1}{n(n-1)m^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n k(X_{ij}, Y_{st}) \right) \tag{32}
\end{aligned}$$

$$\begin{aligned}
& \mathbb{V} \left( \frac{1}{nm^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, Y_{it}) \right) \\
& \stackrel{\text{TV}}{=} \mathbb{V} \left( \mathbb{E} \left[ \frac{1}{nm^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, Y_{it}) \mid P_1 \dots P_n, Q_1 \dots Q_n \right] \right) \\
& \quad + \mathbb{E} \left[ \mathbb{V} \left( \frac{1}{nm^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, Y_{it}) \mid P_1 \dots P_n, Q_1 \dots Q_n \right) \right] \\
& \stackrel{\text{iid}}{=} \mathbb{V} \left( \frac{1}{nm^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m \mathbb{E} [k(X_{ij}, Y_{it}) \mid P_i, Q_i] \right) \\
& \quad + \mathbb{E} \left[ \frac{1}{n^2} \sum_{i=1}^n \mathbb{V} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, Y_{it}) \mid P_i, Q_i \right) \right] \\
& \stackrel{\text{iid}}{=} \mathbb{V} \left( \frac{1}{n} \sum_{i=1}^n \mathbb{E} [k(X_{i1}, Y_{i1}) \mid P_i, Q_i] \right) + \frac{1}{n} \mathbb{E} \left[ \mathbb{V} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{1j}, Y_{1t}) \mid P_1, Q_1 \right) \right] \\
& \stackrel{\text{iid}}{=} \frac{1}{n} \mathbb{V} [\langle P \mid k \mid Q \rangle] + \frac{1}{n} \mathbb{E} \left[ \mathbb{V} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{1j}, Y_{1t}) \mid P_1, Q_1 \right) \right] \\
& \stackrel{\text{iid}}{=} \frac{1}{n} \underbrace{\mathbb{V} [\langle P \mid k \mid Q \rangle]}_{\eta_1 :=} + \frac{1}{nm^2} \underbrace{\mathbb{E} [\mathbb{V} (k(X_{11}, Y_{11}) \mid P_1, Q_1)]}_{\eta_2 :=}
\end{aligned} \tag{33}$$

and for the second term we have

$$\begin{aligned}
& \mathbb{V} \left( \frac{1}{n(n-1)m^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n k(X_{ij}, Y_{st}) \right) \\
& \stackrel{\text{TV}}{=} \mathbb{V} \left( \underbrace{\mathbb{E} \left[ \frac{1}{n(n-1)m^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{t=1}^m k(X_{ij}, Y_{st}) \mid P_1 \dots P_n, Q_1 \dots Q_n \right]}_{\text{(III)} :=} \right) \\
& \quad + \mathbb{E} \left[ \underbrace{\mathbb{V} \left( \frac{1}{n(n-1)m^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{t=1}^m k(X_{ij}, Y_{st}) \mid P_1 \dots P_n, Q_1 \dots Q_n \right)}_{\text{(IV)} :=} \right].
\end{aligned} \tag{34}$$

Due to the length of the expression, we first solve (III) and then (IV).

$$\begin{aligned}
\text{(III)} &\stackrel{\text{iid}}{=} \mathbb{V} \left( \frac{1}{n(n-1)m^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{t=1}^m \mathbb{E}[k(X_{ij}, Y_{st}) | P_i, Q_s] \right) \\
&\stackrel{\text{iid}}{=} \mathbb{V} \left( \frac{1}{n(n-1)m^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{t=1}^m \mathbb{E}[k(X_{i1}, Y_{s1}) | P_i, Q_s] \right) \\
&= \mathbb{V} \left( \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \langle P_i | k | Q_s \rangle \right) \\
&= \mathbb{E} \left[ \left( \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \langle P_i | k | Q_s \rangle \right)^2 \right] - \left( \mathbb{E} \left[ \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \langle P_i | k | Q_s \rangle \right] \right)^2 \\
&\stackrel{\text{iid}}{=} \frac{1}{n^2(n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{j=1}^n \sum_{\substack{t=1 \\ t \neq j}}^n \mathbb{E}[\langle P_i | k | Q_s \rangle \langle P_j | k | Q_t \rangle] - \langle \mathbb{E}[P] | k | \mathbb{E}[Q] \rangle^2 \\
&= \frac{1}{n^2(n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \mathbb{E}[\langle P_i | k | Q_s \rangle^2] + \frac{1}{n^2(n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{t=1 \\ t \neq j \\ t \neq s}}^n \mathbb{E}[\langle P_i | k | Q_s \rangle \langle P_i | k | Q_t \rangle] \\
&\quad + \frac{1}{n^2(n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{j=1 \\ j \neq i \\ j \neq s}}^n \mathbb{E}[\langle P_i | k | Q_s \rangle \langle P_j | k | Q_s \rangle] \\
&\quad + \frac{1}{n^2(n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{j=1 \\ j \neq i \\ j \neq s}}^n \sum_{\substack{t=1 \\ t \neq i \\ t \neq j \\ t \neq s}}^n \mathbb{E}[\langle P_i | k | Q_s \rangle \langle P_j | k | Q_t \rangle] - \langle \mathbb{E}[P] | k | \mathbb{E}[Q] \rangle^2 \\
&\stackrel{\text{iid}}{=} \frac{1}{n(n-1)} \mathbb{E}[\langle P_1 | k | Q_2 \rangle^2] + \frac{n-2}{n(n-1)} \mathbb{E}[\langle P_1 | k | Q_2 \rangle \langle P_1 | k | Q_3 \rangle] \\
&\quad + \frac{n-2}{n(n-1)} \mathbb{E}[\langle P_2 | k | Q_1 \rangle \langle P_3 | k | Q_1 \rangle] + \frac{(n-2)(n-3)}{n(n-1)} \langle \mathbb{E}[P] | k | \mathbb{E}[Q] \rangle^2 - \langle \mathbb{E}[P] | k | \mathbb{E}[Q] \rangle^2 \\
&= \frac{1}{n(n-1)} \mathbb{E}[\langle P_1 | k | Q_2 \rangle^2] + \frac{n-2}{n(n-1)} \mathbb{E}[\langle P_1 | k | Q_2 \rangle \langle P_1 | k | Q_3 \rangle] \\
&\quad + \frac{n-2}{n(n-1)} \mathbb{E}[\langle P_2 | k | Q_1 \rangle \langle P_3 | k | Q_1 \rangle] + \frac{(n-2)(n-3) - n(n-1)}{n(n-1)} \langle \mathbb{E}[P] | k | \mathbb{E}[Q] \rangle^2 \\
&= \frac{1}{n(n-1)} \mathbb{E}[\langle P_1 | k | Q_2 \rangle^2] + \frac{n-2}{n(n-1)} \mathbb{E}[\langle P_1 | k | Q_2 \rangle \langle P_1 | k | Q_3 \rangle] \\
&\quad + \frac{n-2}{n(n-1)} \mathbb{E}[\langle P_2 | k | Q_1 \rangle \langle P_3 | k | Q_1 \rangle] + \frac{(n-2)(n-3) - n(n-1)}{n(n-1)} \langle \mathbb{E}[P] | k | \mathbb{E}[Q] \rangle^2 \\
&= \frac{1}{n(n-1)} \mathbb{E}[\langle P_1 | k | Q_2 \rangle^2] + \frac{n-2}{n(n-1)} \mathbb{E}[\langle P_1 | k | Q_2 \rangle \langle P_1 | k | Q_3 \rangle] \\
&\quad + \frac{n-2}{n(n-1)} \mathbb{E}[\langle P_2 | k | Q_1 \rangle \langle P_3 | k | Q_1 \rangle] - \frac{4n-6}{n(n-1)} \langle \mathbb{E}[P] | k | \mathbb{E}[Q] \rangle^2 \\
&\stackrel{\text{(i)}}{=} \frac{1}{n(n-1)} \eta_3 + \frac{n-2}{n(n-1)} \eta_4
\end{aligned} \tag{35}$$

(i) with  $\eta_3 := \mathbb{E} [\langle P_1 | k | Q_2 \rangle^2] - 2 \langle \mathbb{E}[P] | k | \mathbb{E}[Q] \rangle^2$  and  $\eta_4 := \mathbb{E} [\langle P_1 | k | Q_2 \rangle \langle P_1 | k | Q_3 \rangle] + \mathbb{E} [\langle P_2 | k | Q_1 \rangle \langle P_3 | k | Q_1 \rangle] - 4 \langle \mathbb{E}[P] | k | \mathbb{E}[Q] \rangle^2$ .

Due to analogous reasons as in the distributional variance case, we have

$$\begin{aligned}
(\text{IV}) &= \mathbb{E} \left[ \underbrace{\frac{1}{n^2 (n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \mathbb{V} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, Y_{st}) \mid P_i, Q_s \right)}_{(\text{IVa}):=} \right] \\
&+ \mathbb{E} \left[ \underbrace{\frac{1}{n^2 (n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{b=1 \\ b \neq i \\ b \neq s}}^n \text{Cov} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, Y_{st}), \frac{1}{m^2} \sum_{c=1}^m \sum_{d=1}^m k(X_{ic}, Y_{bd}) \mid P_i, Q_b, Q_s \right)}_{(\text{IVb}):=} \right] \\
&+ \mathbb{E} \left[ \underbrace{\frac{1}{n^2 (n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{a=1 \\ a \neq s \\ a \neq i}}^n \text{Cov} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, Y_{st}), \frac{1}{m^2} \sum_{c=1}^m \sum_{d=1}^m k(X_{ac}, Y_{sd}) \mid P_i, P_a, Q_s \right)}_{(\text{IVc}):=} \right].
\end{aligned} \tag{36}$$

Due to the length of the expressions, we again first look at (IVa), then (IVb), and then (IVc).

In the following equation, we use almost the identical steps as in Equation 24:



$$\begin{aligned}
(\text{IVa}) &= \mathbb{E} \left[ \frac{1}{n^2 (n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \mathbb{V} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, Y_{st}) \mid P_i, Q_s \right) \right] \\
&\stackrel{\text{iid}}{=} \frac{1}{n(n-1)} \mathbb{E} \left[ \mathbb{V} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{1j}, Y_{2t}) \mid P_1, Q_2 \right) \right] \\
&= \frac{1}{n(n-1)} \mathbb{E} \left[ \text{Cov} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{1j}, Y_{2t}), \frac{1}{m^2} \sum_{i=1}^m \sum_{s=1}^m k(X_{1i}, Y_{2s}) \mid P_1, Q_2 \right) \right] \\
&= \frac{1}{n(n-1)m^4} \sum_{j=1}^m \sum_{t=1}^m \sum_{i=1}^m \sum_{s=1}^m \mathbb{E} [\text{Cov}(k(X_{1j}, Y_{2t}), k(X_{1i}, Y_{2s}) \mid P_1, Q_2)] \\
&\stackrel{\text{iid}}{=} \frac{1}{n(n-1)m^4} \sum_{j=1}^m \sum_{t=1}^m \mathbb{E} [\text{Cov}(k(X_{1j}, Y_{2t}), k(X_{1j}, Y_{2t}) \mid P_1, Q_2)] \\
&\quad + \frac{1}{n(n-1)m^4} \sum_{j=1}^m \sum_{t=1}^m \sum_{\substack{i=1 \\ i \neq j}}^m \mathbb{E} [\text{Cov}(k(X_{1j}, Y_{2t}), k(X_{1i}, Y_{2t}) \mid P_1, Q_2)] \\
&\quad + \frac{1}{n(n-1)m^4} \sum_{j=1}^m \sum_{t=1}^m \sum_{\substack{s=1 \\ s \neq t}}^m \mathbb{E} [\text{Cov}(k(X_{1j}, Y_{2t}), k(X_{1j}, Y_{2s}) \mid P_1, Q_2)] \\
&\stackrel{\text{iid}}{=} \frac{1}{n(n-1)m^2} \underbrace{\mathbb{E} [\text{Cov}(k(X_{11}, Y_{21}), k(X_{11}, Y_{21}) \mid P_1, Q_2)]}_{\eta_5 :=} \\
&\quad + \frac{m-1}{n(n-1)m^2} \underbrace{\mathbb{E} [\text{Cov}(k(X_{11}, Y_{21}), k(X_{12}, Y_{21}) \mid P_1, Q_2)]}_{\eta_6 :=} \\
&\quad + \frac{m-1}{n(n-1)m^2} \underbrace{\mathbb{E} [\text{Cov}(k(X_{11}, Y_{21}), k(X_{11}, Y_{22}) \mid P_1, Q_2)]}_{\eta_7 :=}.
\end{aligned} \tag{37}$$

Next, we have analogous to Equation 25

$$\begin{aligned}
(\text{IVb}) &= \mathbb{E} \left[ \frac{1}{n^2 (n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{b=1 \\ b \neq i \\ b \neq s}}^n \text{Cov} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, Y_{st}), \frac{1}{m^2} \sum_{c=1}^m \sum_{d=1}^m k(X_{ic}, Y_{bd}) \mid P_i, Q_s, Q_b \right) \right] \\
&= \mathbb{E} \left[ \frac{1}{n^2 (n-1)^2 m^4} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{b=1 \\ b \neq i \\ b \neq s}}^n \sum_{j=1}^m \sum_{t=1}^m \sum_{d=1}^m \text{Cov} (k(X_{ij}, Y_{st}), k(X_{ij}, Y_{bd}) \mid P_i, Q_s, Q_b) \right] \\
&\stackrel{\text{iid}}{=} \mathbb{E} \left[ \frac{1}{n^2 (n-1)^2 m^4} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{b=1 \\ b \neq i \\ b \neq s}}^n \sum_{j=1}^m \sum_{t=1}^m \sum_{d=1}^m \text{Cov} (k(X_{11}, Y_{21}), k(X_{11}, Y_{31}) \mid P_1, Q_2, Q_3) \right] \\
&= \frac{n-2}{n(n-1)m} \underbrace{\mathbb{E} [\text{Cov} (k(X_{11}, Y_{21}), k(X_{11}, Y_{31}) \mid P_1, Q_2, Q_3)]}_{\eta_9 :=}
\end{aligned} \tag{38}$$

and in an almost identical manner

$$\begin{aligned}
(\text{IVc}) &= \mathbb{E} \left[ \frac{1}{n^2 (n-1)^2} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{a=1 \\ a \neq s \\ a \neq i}}^n \text{Cov} \left( \frac{1}{m^2} \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, Y_{st}), \frac{1}{m^2} \sum_{c=1}^m \sum_{d=1}^m k(X_{ac}, Y_{sd}) \mid P_i, P_a, Q_s \right) \right] \\
&= \mathbb{E} \left[ \frac{1}{n^2 (n-1)^2 m^4} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{b=1 \\ b \neq i \\ b \neq s}}^n \sum_{j=1}^m \sum_{t=1}^m \sum_{d=1}^m \text{Cov} (k(X_{ij}, Y_{st}), k(X_{ac}, Y_{sd}) \mid P_i, P_a, Q_s) \right] \\
&\stackrel{\text{iid}}{=} \mathbb{E} \left[ \frac{1}{n^2 (n-1)^2 m^4} \sum_{i=1}^n \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{\substack{b=1 \\ b \neq i \\ b \neq s}}^n \sum_{j=1}^m \sum_{t=1}^m \sum_{d=1}^m \text{Cov} (k(X_{11}, Y_{31}), k(X_{21}, Y_{31}) \mid P_1, P_2, Q_3) \right] \\
&= \frac{n-2}{n(n-1)m} \underbrace{\mathbb{E} [\text{Cov} (k(X_{11}, Y_{31}), k(X_{21}, Y_{31}) \mid P_1, P_2, Q_3)]}_{\eta_{10} :=}.
\end{aligned} \tag{39}$$

The only term left is

$$\begin{aligned}
& \text{Cov} \left( \frac{1}{nm^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m k(X_{ij}, Y_{it}), \frac{1}{n(n-1)m^2} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n k(X_{ij}, Y_{st}) \right) \\
&= \frac{1}{n^2(n-1)m^4} \sum_{i=1}^n \sum_{j=1}^m \sum_{t=1}^m \sum_{\substack{o=1 \\ s \neq o}}^n \sum_{p=1}^m \sum_{\substack{s=1 \\ s \neq o}}^n \sum_{r=1}^m \text{Cov}(k(X_{ij}, Y_{it}), k(X_{op}, Y_{sr})) \\
&\stackrel{(i)}{=} \frac{1}{n^2(n-1)m^4} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \sum_{p=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{r=1}^m \text{Cov}(k(X_{ij}, Y_{it}), k(X_{ip}, Y_{sr})) \\
&\quad + \frac{1}{n^2(n-1)m^4} \sum_{i=1}^n \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \sum_{p=1}^m \sum_{\substack{s=1 \\ s \neq i}}^n \sum_{r=1}^m \text{Cov}(k(X_{ij}, Y_{it}), k(X_{sp}, Y_{ir})) \\
&\stackrel{\text{iid}}{=} \frac{1}{nm^3} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \sum_{p=1}^m \text{Cov}(k(X_{1j}, Y_{1t}), k(X_{1p}, Y_{21})) + \frac{1}{nm^3} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \sum_{r=1}^m \text{Cov}(k(X_{1j}, Y_{1t}), k(X_{21}, Y_{1r})) \\
&= \frac{1}{nm^3} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \text{Cov}(k(X_{1j}, Y_{1t}), k(X_{1j}, Y_{21})) + \frac{1}{nm^3} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \text{Cov}(k(X_{1j}, Y_{1t}), k(X_{1t}, Y_{21})) \\
&\quad + \frac{1}{nm^3} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \sum_{\substack{p=1 \\ p \neq j \\ p \neq t}}^m \text{Cov}(k(X_{1j}, Y_{1t}), k(X_{1p}, Y_{21})) + \frac{1}{nm^3} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \text{Cov}(k(X_{1j}, Y_{1t}), k(X_{21}, Y_{1j})) \\
&\quad + \frac{1}{nm^3} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \text{Cov}(k(X_{1j}, Y_{1t}), k(X_{21}, Y_{1t})) + \frac{1}{nm^3} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \sum_{\substack{p=1 \\ p \neq j \\ p \neq t}}^m \text{Cov}(k(X_{1j}, Y_{1t}), k(X_{21}, Y_{1p})) \\
&\stackrel{\text{iid}}{=} \frac{1}{nm^3} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \text{Cov}(k(X_{11}, Y_{12}), k(X_{11}, Y_{21})) + \frac{1}{nm^3} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \text{Cov}(k(X_{11}, Y_{12}), k(X_{12}, Y_{21})) \\
&\quad + \frac{1}{nm^3} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \sum_{\substack{p=1 \\ p \neq j \\ p \neq t}}^m \text{Cov}(k(X_{11}, Y_{12}), k(X_{13}, Y_{21})) + \frac{1}{nm^3} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \text{Cov}(k(X_{11}, Y_{12}), k(X_{21}, Y_{11})) \\
&\quad + \frac{1}{nm^3} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \text{Cov}(k(X_{11}, Y_{12}), k(X_{21}, Y_{12})) + \frac{1}{nm^3} \sum_{j=1}^m \sum_{\substack{t=1 \\ t \neq j}}^m \sum_{\substack{p=1 \\ p \neq j \\ p \neq t}}^m \text{Cov}(k(X_{11}, Y_{12}), k(X_{21}, Y_{13})) \\
&= \frac{m-1}{nm^2} \underbrace{(\text{Cov}(k(X_{11}, Y_{12}), k(X_{11}, Y_{21})) + \text{Cov}(k(X_{11}, Y_{12}), k(X_{21}, Y_{11})))}_{\eta_{11}:=} \\
&\quad + \frac{m-1}{nm^2} \underbrace{(\text{Cov}(k(X_{11}, Y_{12}), k(X_{21}, Y_{12})) + \text{Cov}(k(X_{11}, Y_{12}), k(X_{12}, Y_{21})))}_{\eta_{12}:=} \\
&\quad + \frac{(m-1)(m-2)}{nm^2} \underbrace{(\text{Cov}(k(X_{11}, Y_{12}), k(X_{13}, Y_{21})) + \text{Cov}(k(X_{11}, Y_{12}), k(X_{21}, Y_{13})))}_{\eta_{13}:=}.
\end{aligned} \tag{40}$$

By combining all previous equations, we get

$$\begin{aligned}
& \mathbb{V} \left( \widehat{\text{Cov}}_k^{(n,m)}(\mathbf{X}, \mathbf{Y}) \right) \\
&= \underbrace{\frac{1}{n} \eta_1 + \frac{n-2}{n(n-1)} \eta_4 + \frac{-2(m-1)(m-2)}{nm^2} \eta_{13}}_{\mathcal{O}\left(\frac{1}{n}\right)} + \underbrace{\frac{1}{n(n-1)} \eta_3}_{\mathcal{O}\left(\frac{1}{n^2}\right)} \\
&\quad + \underbrace{\frac{-2(m-1)}{nm^2} (\eta_{11} + \eta_{12}) + \frac{n-2}{n(n-1)m} (\eta_9 + \eta_{10})}_{\mathcal{O}\left(\frac{1}{nm}\right)} \\
&\quad + \underbrace{\frac{1}{nm^2} \eta_2}_{\mathcal{O}\left(\frac{1}{nm^2}\right)} + \underbrace{\frac{m-1}{n(n-1)m^2} (\eta_6 + \eta_7)}_{\mathcal{O}\left(\frac{1}{n^2m}\right)} + \underbrace{\frac{1}{n(n-1)m^2} \eta_5}_{\mathcal{O}\left(\frac{1}{n^2m^2}\right)}.
\end{aligned} \tag{41}$$