

ANALYSIS OF RADIO LOCALISER NETWORKS UNDER DISTRIBUTION SHIFT

Anonymous authors

Paper under double-blind review

ABSTRACT

Deploying radio frequency (RF) localisation systems invariably entails non-trivial effort, particularly for the latest learning-based breeds. There has been little prior work on characterising and comparing how learnt localiser networks can be deployed in the field under real-world RF distribution shifts. In this paper, we present RadioBench: a suite of 8 learnt localiser nets from the state-of-the-art to study and benchmark their real-world deployability, utilising five novel industry-grade datasets. We train 10k models to analyse the inner workings of these learnt localiser nets and uncover their differing behaviours across three performance axes: (i) learning, (ii) proneness to distribution shift, and (iii) localisation. We use insights gained from this analysis to recommend best practices for the deployability of learning-based RF localisation under practical constraints.

1 INTRODUCTION

Decades of radio frequency (RF) localisation research have given us a variety of classic methods (Patwari et al., 2005; Gezici et al., 2005). Newer machine learning incarnations can enhance location estimation considerably (Zanjani et al., 2022; Karmanov et al., 2021), albeit at the expense of proneness to distributional shift in wireless signals. For example, models trained on signals from a warehouse environment may not work well in another different environment (Arnold et al., 2018). If learnt localiser networks are to be productised and deployed, it is imperative that we robustify them. To achieve real-world robustness, we need to understand (i) the performance nuances of learnt localisation models, (ii) when, how, and why do such models work, and (iii) when do they fail.

Robustness to distribution shift (i.e., out of distribution (OOD) generalisation) is an established line of enquiry in mainstream machine learning (Gulrajani & Lopez-Paz, 2020; Hendrycks et al., 2021; Koh et al., 2021). However, there is little in the way of robustness investigations for learnt RF localisation. Though lower dimensional than images, wireless signals are prone to acute variabilities stemming from environment- and/or system-dependent propagation conditions (Tse & Viswanath, 2005), which are hard to control for. Sidestepping this complexity, recent works have incorporated environment-dependent priors (e.g., floorplans) in order to achieve robust learnt RF localisation in that environment (Karmanov et al., 2021; Zanjani et al., 2022; Ghazvinian Zanjani et al., 2021).

In this paper, we seek to understand the practical deployability of learnt RF localiser nets from first principles and without invoking extra robustifying priors. To this end, we build RadioBench: a suite of RF localiser nets from the state-of-the-art. We conduct a systematic comparative study on these localiser nets, utilising five novel industry-grade datasets. We analyse the inner workings of these localiser nets and uncover their differing behaviours across three performance axes: (i) learning, (ii) proneness to distribution shift, and (iii) localisation. Our contributions are:

- We introduce RadioBench: a benchmarking suite of RF localiser nets from the state-of-the-art, as well as a best-in-class classical probabilistic approach.
- We introduce 5 large-scale, industry-grade RF localisation datasets with differing characteristics that pertain to the study of wireless OOD robustness.
- We characterise and contrast the performance of 8 RF localiser methods, training in excess of 10k models in the process. These model configurations span: architecture, representation learning, and domain adaptation methods.

- We find that representation learning and pretraining are most important for OOD robustness in a new RF environment, and that variants based on an autoencoder architecture are the best all-rounder models.

2 PRIMER ON RF LOCALISATION

We consider a system of M synchronised locators that listen for user devices, where each m th locator has known 3D position vector $\mathbf{u}_m = [x_m, y_m, z_m]$, and 3D 3×3 orientation matrix $\mathbf{\Omega}_m$. Let \mathbf{A} be the angle of arrival (AoA) matrix, r the range calculated using time of arrival (ToA) and the speed of the light, then the position of a user device w.r.t. m th locator

$$\mathbf{p}_m = \mathbf{\Omega}_m \mathbf{A} r + \mathbf{u}_m \quad (1)$$

Because user devices and locators are not synchronised, range estimates are biased. This can be compensated by using one locator as reference using time difference of arrival (TDoA). Typically, modern RF localisation relies on estimating the aforementioned two wireless propagation properties ToA and AoA, which together are abbreviated (TAoA).

Challenge in rich scattering. Considering a wireless channel between two radio transceivers, the baseband model of the channel impulse response is given by (Tse & Viswanath, 2005)

$$h(k) = \sum_{p=1}^P \sum_{\ell=0}^{L-1} a_{p,\ell} e^{j(2\pi f_c \tau_p + \phi_{p,\ell})} \text{sinc}\left(k - \frac{\tau_{p,\ell}}{T_s}\right), \quad k = 0, \dots, O-1$$

where $a_p \in \mathbb{R}_+$, $\phi_{p,\ell} \in \mathbb{R}$, $\tau_p \in \mathbb{R}_+$ are respectively the attenuation, phase, and propagation delay of the p th path and ℓ th path cluster. Also $\text{sinc}(x) = \frac{\sin(\pi x)}{\pi x}$ is the normalised sinc function, k is the discrete sampling time, and $O-1$ is the channel order.

It is generally infeasible to estimate the above parameters because they are underdetermined in practical implementations. This is further compounded by environments with rich scattering (i.e., large P and L).

Upper bound. Eq. 1 shows that the best performance can be theoretically achieved using perfect TAoA labels as input to a deep neural net. TAoA, however, are infeasible to measure as groundtruth per deployed environment because it would entail extensive and very expensive surveying campaigns. Deployment surveys typically leverage laser measurements and tens of hours of calibration (Scott & Hazas, 2003). Further, moving from a local coordinate system (i.e., per locator) to a global coordinate system for the environment requires models of that environment and the locator hardware. Therefore, we designate a TAoA-based localiser net as an upper bound on performance that is impractical to implement in the real-world under realistic deployment cost and overhead constraints.

3 MODEL VARIANTS

RadioBench suite compiles all RF localiser net architectures reported in literature. While all facilitate location estimation, these architectures operate on differing input formats, produce differing output formats, as well as deviate in their training details. We believe RadioBench to be the first effort to comprehensively catalogue and evaluate RF localiser nets in order to concretely establish and contrast their performances. Appendix A reviews RF localisation fundamentals and treats learnt localiser net variants in more detail.

3.1 ARCHITECTURES

We evaluate four classes of RF localiser nets: supervised CNN (Chen et al., 2017; Arnold et al., 2019), supervised residual net (ResNet) akin to vision ResNet (He et al., 2016), unsupervised AutoEncoder (AE) (Liu et al., 2018), and unsupervised channel charting (CC) (Studer et al., 2018).

3.2 INPUT-OUTPUT (IO) FORMATS

The above architectures can ingest various representations of input wireless signals. These are: (i) Channel state information (CSI) is the raw measurements obtained from transceiver chips, (ii)

periodograms (PER) is CSI’s 2D Fourier projection, (iii) a feature reduced version of (i) or (ii), and (iv) TAOA are the physical propagation primitives that implicitly encode location (cf., Eq. 1), and are obtained via surveying the environment as discussed in Sec. 2.

The above architectures can also produce multiple output representations. These representations either encode location directly, or encapsulate it indirectly. Specifically, output can be: (i) position estimate, (ii) TAOA primitives, or (iii) latent space that implicitly contains the location intrinsic space.

Tab. 1 lists all valid architecture-IO configurations supported in RadioBench. Specifically for each method, Tab. 1 shows the effective mapping and its optimisation objective, which is minimisation for AE and maximisation for CC. For further details around these methods, consult original literature.

Table 1: RF localiser nets and their valid architecture, input-output configurations, and training objective implemented in RadioBench.

| Detail | Supervised | Autoencoder (AE) | Channel chart (CC) |
|---------------|--|--|---|
| Mapping | $\mathbb{C}^M \rightarrow \mathbb{R}^3$ | $\mathbb{C}^M \rightarrow \mathbb{R}^{M'}$ | $\mathbb{C}^M \rightarrow \mathbb{R}^2$ |
| Optimisation | $\ \mathbf{p}_n - g(\mathbf{f}_n)\ _2^2$ | $\ \mathbf{f}_n - g^{-1}(g(\mathbf{f}_n))\ _2^2$ | $\ d(g(T_a), g(T_p)) - d(g(T_a), g(T_n))\ $ |
| Type | ResNet | CNN | CNN |
| Input | CSI/PER | CSI/PER | $f(\text{CSI/PER})$ |
| Output | Position/TAOA | M' Features | Channel Chart |
| Configuration | CSI2Pos, PER2Pos CSI2TAOA, PER2TAOA | CSI AE, PER AE | CSI CC, PER CC |

3.3 CLASSICAL BASELINE

A best-in-class probabilistic method is also used for benchmarking (Henninger et al., 2022), which we designate as Classical. Classical uses super-resolution techniques to estimate TAOA from CSI. These TAOAs are inputted to a maximum likelihood estimation (MLE) pipeline. Note that Classical results presented throughout paper are averaged per grid position for best-case analysis.

4 FRAMEWORK FOR EMPIRICAL OOD ROBUSTNESS ANALYSIS

We introduce five large-scale, industry-grade datasets that enable us to empirically study the OOD robustness of wireless localiser nets. We review the distribution shift mechanisms at play in these datasets. We then discuss the distribution shift mitigation strategies we deem applicable to the model variants of Sec. 3.

4.1 DATASETS

We utilise a set of empirical industry-grade datasets to study the nuances of radio localiser net variants. We summarise the setup and geometric configurations under which the radio measurement campaign was conducted.

Fig. 1 depicts three physically distinct environments. Within each, six locators listen to mobile users. Each locator is equipped with a 3×3 antenna array. Locators are tightly synchronised using White Rabbit standard (Eidson et al., 2002). Mobile user devices regularly transmit pilot data known to the locators. The locators receive user pilots and estimate their CSI. Fig. 1a shows 3 Arena 1 measurements in blue, green, and red. These correspond to three data collection iterations. Arena 1 and 2 (blue and green) cover the same area but are different due to hardware effects. Arena 3 (red rectangle) corresponds to high-speed driving to simulate a dynamic environment. Fig. 1b & c

Table 2: Datasets utilised in evaluation and their configurations.

| # | Dataset | points | f_c (GHz) | Bandwidth (MHz) | Subcarriers | Antennae | Locators | Groundtruth ^{*†} | Area (m^2) |
|---|------------|--------|-------------|-----------------|-------------|----------|----------|---------------------------|----------------|
| 1 | Arena 1 | 52991 | 3.75 | 100 | 1630 | 8 | 6 | SLAM | 134.532 |
| 2 | Arena 2 | 46266 | 3.75 | 100 | 1630 | 8 | 6 | SLAM | 134.709 |
| 3 | Arena 3 | 2181 | 3.75 | 100 | 1630 | 8 | 6 | SLAM | 57.188 |
| 4 | Industry 1 | 7990 | 3.75 | 100 | 1630 | 8 | 6 | Tachy | 387.872 |
| 5 | Industry 2 | 5037 | 3.75 | 100 | 1630 | 8 | 6 | Tachy | 103.469 |

^{*}SLAM: obtained from simultaneous localization and mapping of mobile robot equipped with Lidar

[†]Tachy: high-precision laser surveying device

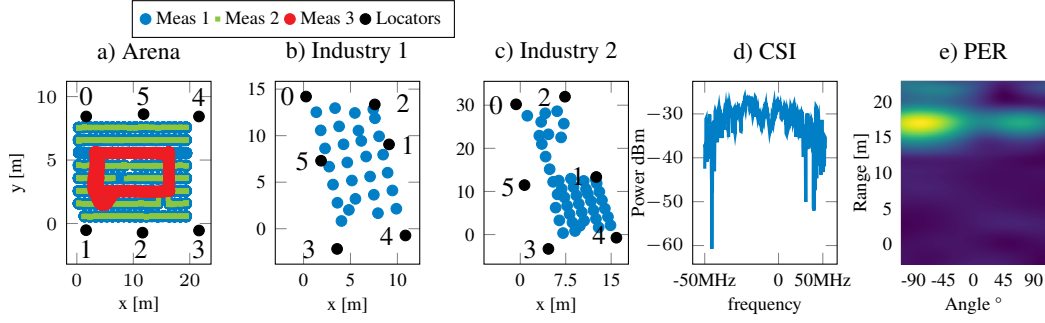


Figure 1: Measurement environments. a) Arena has three different measurement iterations, b) Typical Industrial environment, and c) Harsher industrial environment with rich scatterers. Data examples: d) CSI and e) PER.

correspond to two other industrial environments, with Industry 2 being particularly rich in scattering effects. Fig. 1d & e depict two examples of the input formats discussed in Tab. 1.

Tab. 2 summarises the configurations of our 5 novel datasets.

4.2 DISTRIBUTION SHIFT MECHANISMS

From the 5 datasets listed in Tab. 2, we highlight the following mechanisms that result in distributional shift in RF signals.

(1) Macro environment-induced. Each of the three environments depicted in Fig. 1 comes with its signature set of propagation conditions. These propagation conditions are largely a function of the geometry of the environment as well as the spatial configuration of reflective surfaces present within, e.g., metallic machinery, furniture, partition walls and their material composition, etc. We designate environmental signatures as macro-level effects that shift the bulk of the distribution of RF signals.

(2) Micro locator-induced. The 3D position and orientation of locators within the environment affect how they measure the statistics of user RF signals. That is, a locator will also modulate the distribution of the RF signals it receives. We designate locator signatures as micro-level effects that further shift the distribution of RF signals.

(3) Micro scattering-induced. Dynamic activities within the environment induce scattering effects that modulate the distribution of the RF signals. Example scatterers include moving robots and people. We designate scattering as micro-level effects that further shift the distribution of RF signals.

(4) Misc. For completeness, there are multiple other factors that affect the distribution of RF signals. Examples include hardware- and frequency-dependent effects. We, however, are mainly interested in shift mechanisms 1-3 in this work as captured by our 5 empirical datasets.

4.3 DISTRIBUTION SHIFT MITIGATIONS

There are a wide range of methods from the state-of-the-art that enhances robustness and generalisation on unseen distribution shifts. However, the relative performance of these methods varies largely across modalities, datasets, and distribution shifts (Hendrycks et al., 2021; Koh et al., 2021; Wiles et al., 2021). Further, there is little prior experience in adapting some of these concepts to the RF localiser net setting we study herein.

Loss landscape. Not all models are created equal. For all models described in Sec. 3, we visualise a dataset’s loss landscape using (Li et al., 2018). This analysis is motivated by the observation that the landscape geometry affects generalisation dramatically (Li et al., 2018). All things being equal, we would therefore favour model variants that exhibit flatter loss landscape geometry. Our intuition is that a flatter loss landscape would readily support a weak form of generalisability.

Fine-tuning. Fine-tuning is a consistent indicator of the quality of zero-shot models (Radford et al., 2021; Wortsman et al., 2022). It is decoupled from some modality-specific recipes such as data augmentation. Therefore, we employ simple universal fine-tuning for zero-shot model variants from Sec. 3 to gauge their relative robustness and generalisability to unseen distribution shifts.

4.4 LEARNABILITY CONDITIONS

We characterise various aspects around the learnability of the model variants described in Sec. 3.

Active label density. We investigate the required number of labels for validation loss convergence. We employ active learning strategies to glean comparative insights on the learning behaviour of localiser model variants.

Latent space. Some model variants utilise a latent space that implicitly encodes location. We investigate the resultant shift in the latent space as a function of macro and micro RF signal distribution shifts.

Regression head protocol. For model variants with a latent space, we investigate the feasibility of a regressor head on top of a frozen backbone that is trained on a different dataset. We intuit that if quality features have been learnt, their projection would still perform competitively w.r.t. regressing location estimates notwithstanding distribution shift.

5 EXPERIMENTS

We evaluate 8 RF localiser nets. We conduct a comprehensive analysis to quantify performance aspects around: (i) learnability, (ii) proneness to distribution shift, and (iii) localisation. We use our 5 industry-grade datasets (cf., Tab. 2) for all experiments. We either use all or a subset of localiser model variants (cf., Tab. 1) depending on suitability for a given experiment. We begin by distilling our experimental findings into a concrete set of key takeaways.

As discussed in Sec. 2, method TAoA2Pos in all analyses represents an upper bound on performance. This is because in practical deployments, surveying groundtruth TAoAs is prohibitively expensive.

5.1 TAKEAWAYS

1 – Learnability: Under smart sample selection criterion, training samples of the order of the spatial grid suffice for convergence. We observe that models CSI AE, CSI2TAoA, and PER2TAoA converge after selecting a number samples (via active learning) comparable to the number of spatial grid locations. Inspecting Fig. 2, this happens after around 2.7k samples. This is inline with Arena 1 dataset that has around 53k data points of which 2.7k are spatial grid locations (cf., Fig. 1a).

2 – Learnability: Model variants exhibit qualitative learning differences. Methods that directly map to position (i.e., CSI2Pos and PER2Pos) are poor learners. This phenomenon is most evident in Fig. 2’s entropy and margin sampling losses where CSI2Pos shows a sudden drop around 2.7k




Figure 2: Active learning (AL) for model variants on Arena 1. AL criteria help models converge faster in required training samples. Required number of training samples is of the order of a dataset’s spatial location sampling grid. Some model variants are qualitatively better than others from a learning standpoint.

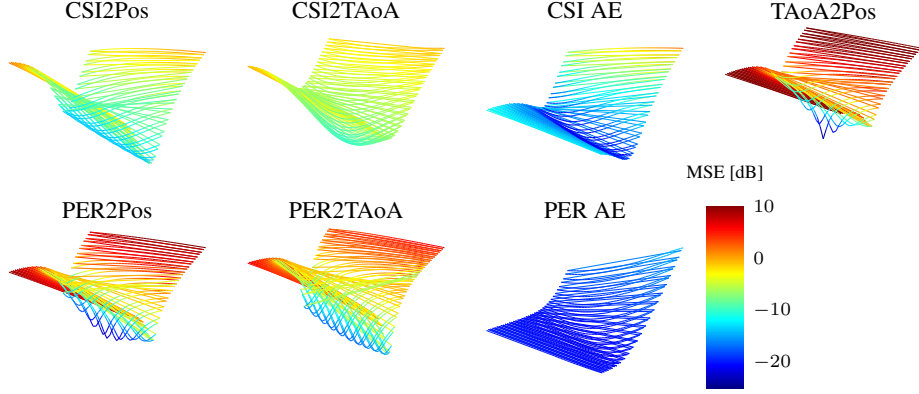


Figure 3: Loss landscapes of model variants on Arena 1. Loss landscapes are a proxy to inspecting the convex hull of model variants. The geometry of a model’s loss landscape is indicative of its generalisability.

samples, which hints at memorisation (i.e., fingerprinting of environment). In contrast, methods that indirectly map to a local ambient space (i.e., CSI AE, CSI2TAoA, and PER2TAoA) do not exhibit such a waterfall effect in their learning loss.

3 – Shift: Model architecture and training details imbue a weak sense of generalisability by construction. RF localiser nets based on an AE architecture (i.e., CSI & PER AEs in Fig. 3) has the best zero-shot weak generalisability owing to flatter loss landscape (Li et al., 2018) (cf., Sec. 4.3). CSI & PER AEs are agnostic to any position-dependent information (direct or indirect) and are trained only on a reconstruction loss. In contrast, methods that use information that encodes position during training (i.e., TAoA2Pos and Per2Pos) are at the opposite end of convexity steepness, necessitating more work for shift mitigation.

Table 3: Zero-shot OOD performance. Performance reported in terms of median error for position and/or azimuth/elevation/range.

| Input | CSI | PER | TAoA | CSI | PER |
|----------------------|----------|----------|----------|--------------------|---------------------|
| Output | Position | Position | Position | TAoA | TAoA |
| Train → Test | | | | | |
| Arena 1 → Arena 1 | 0.07 m | 0.03 m | 0.03 m | 0.03°/0.02°/0.01 m | 0.05°/0.02°/0.007 m |
| Arena 1 → Arena 2 | 0.65 m | 0.19 m | 0.03 m | 0.31°/0.12°/0.04 m | 0.05°/0.02°/0.007 m |
| Arena 1 → Arena 3 | 1.09 m | 0.62 m | 0.11 m | 0.75°/0.37°/0.09 m | 0.17°/0.07°/0.03 m |
| Arena 1 → Industry 1 | 4.70 m | 4.31 m | 2.76 m | 5.73°/1.46°/0.35 m | 0.49°/0.18°/0.07 m |
| Arena 1 → Industry 2 | 8.44 m | 7.83 m | 6.01 m | 3.49°/1.89°/0.97 m | 5.88°/1.54°/0.50 m |

4 – Shift: Zero-shot OOD performance corroborates that not all models are created equal w.r.t. robustness. The representation used for the mapping between the input and output within RF localiser nets has a large bearing on the network’s zero-shot OOD robustness. To examine this, we train all model variants on Arena 1 and test them on all datasets. We conduct one-off calibration of the models’ mean position and TAoA estimates to coarsely correct for distribution shifts. Tab. 3 summarises the median performance obtained under this zero-shot domain adaptation setting. We can see that models with direct position mapping significantly underperform their TAoA mapping counterparts. This finding is another restatement of the observation in Takeaway 3 illustrated in Fig. 3.

5 – Shift: Fine-tuning on few hundred labels achieves good domain adaptation performance across macro and micro distribution shifts. Simple fine-tuning on a new domain is a robust means to significant performance boosting irrespective of the shift mechanism nature and magnitude. Inspecting Fig. 4, a fraction of the label density of the spatial grid is required, i.e., around 1/10th to 1/5th depending upon dataset and model variant. AE variants (i.e., CSI & PER AEs) are most effective. When fine-tuned, pretrained localiser nets outperform the best classic baseline by at least $2\times$ to $3\times$ at the 50th %ile error, again depending upon dataset and model variant.





Figure 4: Performance of fine-tuned pretrained models.

Figure 5: Wasserstein (\mathcal{W}) distances matrix that quantifies the amount of work needed to translate Arena 1 to other macro and micro domains for the CSI AE latent space. \mathcal{W} is averaged across the spatial sampling grid.

6 – Shift: The latent space mirrors macro and micro distribution shifts. The high dimensionality of the AE latent space affords avenues for applying wider range of domain generalisation techniques. To quantify the inter-dataset shifts seen in AE latent space, Fig. 5 measures the Wasserstein distances (Peyré et al., 2017) across our 5 datasets—relative to a model trained on Arena 1. Locator-induced micro shifts cluster the latent spaces of Arena 1 \rightarrow Arena 2-3 into two groups: 0-2 and 3-5. This is in line with the physical nature of Arena configurations as groups 0-2 and 3-5 differ in their orientation for maximum spatial coverage (see Fig. 1a). On Arena 1 \rightarrow Industry 1-2, the picture is more nuanced. For instance, Arena 1 \rightarrow Industry 2 shows that locator 0, which is especially rich in scattering, exhibits larger distances to all other locators.




Figure 6: Zero-shot performance of channel charting. CT & TW are two metrics to quantify the goodness of produced CCs (Studer et al., 2018). (CT, TW) tuples respectively from left to right: (0.98, 0.96), (0.94, 0.93), (0.97, 0.94), (0.89, 0.91), (0.96, 0.92). Colourmap denotes time evolution.

7 – Shift: Distribution shifts make the zero-shot performance of channel charting meaningless. Channel charting (CC) has two drawbacks. First, CC outputs low dimensional mapping of the environment (2D, cf. Tab. 1) that struggles to faithfully embed the global structure of CSI (Karmanov et al., 2021). (Karmanov et al., 2021) addressed this dimensionality bottleneck by inflating CC to be more AE-like, which makes us question the use of CC (and not AE) in the first place. Second, CC requires continuity in sampling the physical environment in order to enforce spatial coherence in its 2D mapping. Together, these two drawbacks severely hamper the zero-shot application of CC on unseen domains. Concretely, Fig. 6 depicts two CCs trained for Arena 1 and Arena 3. Referring to the physical space Arena 1 and Arena 3 traverse in Fig. 1a, we note two trajectories: blue for Arena 1 and red for Arena 3. It is interesting to note that when feeding Arena 3 data into Arena 1 CC (Arena 1 \rightarrow Arena 3 in Fig. 6), we obtain a CC that largely overlaps with Arena 1 CC despite the completely different physical trajectories in Fig. 1a (i.e., blue vs. red). This exposes a major limitation of CC w.r.t. zero-shot performance on unseen data.

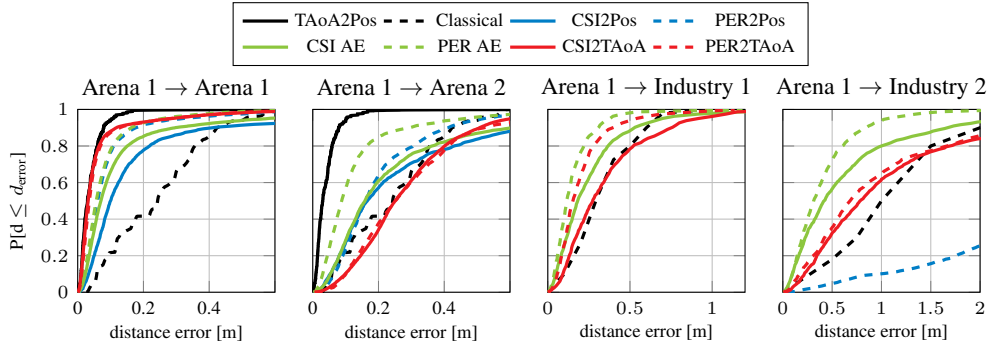


Figure 7: Localisation perf. with pretraining on Arena 1 transferred to other domains. Variants xAE and x2TAoA use the frozen pretrained backbones and nonlinear heads trained on 1k samples. The monolithic x2Pos variants use full finetuning on 1k samples.

8 – Localisation: Pretraining is a powerful tool for building environment-specific localiser nets.

The combination of a pretrained backbone and a specialised nonlinear head is a powerful formula for easier deployment of localisation services. We pretrain backbones for the AE and TAoA configurations on Arena 1. With 1k samples, we then train nonlinear heads on top of these frozen backbones to specialise for other datasets. Results in Fig. 7 show that pretraining indeed gives strong performance on other unseen domains. Specifically, the PER AE configuration outperforms classical by a large margin, while CSI AE, CSI2TAoA, PER2TAoA exhibit competitive performances. Classical struggles particularly in scattering-rich Industry 2. Note that in Arena 1 \rightarrow Arena 1 under no distribution shift, all variants outperform Classical. It is only under distribution shift that we observe clear differentiation in the robustness of the pretrained backbones of model variants.

5.2 PRACTICAL TIPS

Based on our analyses, we would recommend the following tips.

1 – Representation learning is important. Multiple analyses point to qualitative and quantitative differentiation between model variants from RadioBench. Specifically, moving directly to position and not the latent space results in poor OOD robustness. Using a high-dimensional latent space improves robustness to macro and micro distribution shifts, e.g., environment-induced and rich scattering.

2 – Pretraining is powerful. Uniformly across analyses, we found pretraining to be effective for enhancing OOD robustness. Further, fine-tuning pretrained representations seems to only require a small fraction of labels (relative to spatial sampling grid size) for good domain adaptation in RF localiser nets.

3 – The increased dimensionality of latent space helps. Using a latent space with enough dimensionality enhances robustness to distribution shift as well as boosts accuracy.

6 RELATED WORK

Learning-based localisation. Newer localisation techniques use machine learning. The rationale is that data-driven learning is able to model and compensate for the sources of error that limit the performance of classic methods. Machine learning can be applied to both radar techniques (Zhao et al., 2018b;a), or device-based methods (Arnold et al., 2019; Studer et al., 2018; Khatab et al., 2017). Learning-based localisation can either be: supervised (Decurninge et al., 2018; Arnold et al., 2019), or unsupervised (Studer et al., 2018).

Related works exploit prior floorplan information and feature learning to achieve robustness in RF localisation (Ghazvinian Zanjani et al., 2021; Gill et al., 2021; Kadambi et al., 2022; Karmanov et al., 2021). Work in (Zanjani et al., 2022) demonstrates that adding a large metal reflector in an environment causes significant distribution shift that hampers learnt localisation performance. In contrast to this prior art, our work is the first to: (1) investigate robustness to distribution shift

on empirical large-scale data, (2) elucidate how distribution shift impacts learnt localiser variants non-uniformly, and (3) recommend best practices that help enhance robustness.

Benchmarking out of distribution (OOD) robustness. Multiple related works study and benchmark OOD robustness and generalisation as they pertain to mainstream modalities in machine learning (Hendrycks et al., 2021; Koh et al., 2021; Wiles et al., 2021). Informed and inspired by these works, in this paper we contribute analyses that: (1) compile all RF localiser net variants from prior art, and (2) benchmark their performance from an OOD robustness standpoint. The latter is crucial for the maturity and deployment of such models in a particularly distribution shift-prone RF environment.

7 LIMITATIONS

We took first steps towards studying the robustness barrier to real-world deployability of learnt RF localisation. Robustness to OOD is the number 1 issue outstanding in prior art. Naturally, there are many more avenues of future investigations. A few come to mind: (1) broader cross-fertilisation of generalisation methods from mainstream modalities for RF localisation, (2) investigating hardware-induced distribution shift, e.g., antenna configurations, (3) smarter selection criteria of samples to aid distribution shift mitigation, and (4) incorporating sampling in time to further aid robustness.

8 CONCLUSION

In this work, we build a comprehensive framework to analyse and benchmark RF localiser net variants. We introduce 5 novel and large-scale datasets curated in industrial premises. The 5 datasets make possible the empirical study of learning-based RF localisation, and its robustness to macro- and micro-induced distribution shift effects. Our characterisation shows that localiser net variants without a latent space struggle under distributional shift. Our characterisation also shows that there exists a trade-off between accuracy and robustness (including against rich scattering), where AE variants seem to meet this trade-off and perform well. We distil our findings into a set of concrete takeaways, a number of practical tips, and open research directions. We hope that our benchmarking framework would help foster future research towards realising accurate, rapid, and robust deployments of learnt RF localisation.

REFERENCES

- Maximilian Arnold, Sebastian Doerner, Sebastian Cammerer, and Stephan Ten Brink. On deep learning-based massive mimo indoor user localization. In *2018 IEEE 19th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pp. 1–5, 2018.
- Maximilian Arnold, Jakob Hoydis, and Stephan ten Brink. Novel massive mimo channel sounding data applied to deep learning-based indoor positioning. In *SCC 2019; 12th International ITG Conference on Systems, Communications and Coding*, pp. 1–6. VDE, 2019.
- H. Chen, Y. Zhang, W. Li, X. Tao, and P. Zhang. ConFi: Convolutional Neural Networks Based Indoor Wi-Fi Localization Using Channel State Information. *IEEE Access*, 5:18066–18074, 2017. doi: 10.1109/ACCESS.2017.2749516.
- Alexis Decurninge, Luis García Ordóñez, Paul Ferrand, He Gaoning, Li Bojie, Zhang Wei, and Maxime Guillaud. Csi-based outdoor localization for massive mimo: Experiments with a learning approach. In *2018 15th International Symposium on Wireless Communication Systems (ISWCS)*, pp. 1–6. IEEE, 2018.
- John C Eidson, Mike Fischer, and Joe White. Ieee-1588 standard for a precision clock synchronization protocol for networked measurement and control systems. In *Proceedings of the 34th Annual Precise Time and Time Interval Systems and Applications Meeting*, pp. 243–254, 2002.
- Sinan Gezici, Zhi Tian, Georgios B Giannakis, Hisashi Kobayashi, Andreas F Molisch, H Vincent Poor, and Zafer Sahinoglu. Localization via ultra-wideband radios: a look at positioning aspects for future sensor networks. *IEEE signal processing magazine*, 22(4):70–84, 2005.

- Farhad Ghazvinian Zanjani, Ilia Karmanov, Hanno Ackermann, Daniel Dijkman, Simone Merlin, Max Welling, and Fatih Porikli. Modality-agnostic topology aware localization. *Advances in Neural Information Processing Systems*, 34:10457–10468, 2021.
- Kuldeep S. Gill, Son Nguyen, Myo M. Thein, and Alexander M. Wyglinski. Three-way deep neural network for radio frequency map generation and source localization. *CoRR*, abs/2111.12175, 2021.
- Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. *arXiv preprint arXiv:2007.01434*, 2020.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8340–8349, 2021.
- Marcus Henninger, Traian E Abrudan, Silvio Mandelli, Maximilian Arnold, Stephan Saur, Veli-Matti Kolmonen, Siegfried Klein, Thomas Schlitter, and Stephan Ten Brink. Probabilistic 5g indoor positioning proof of concept with outlier rejection. In *2022 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*, pp. 249–254. IEEE, 2022.
- Shreya Kadambi, Arash Behboodi, Joseph B. Soriaga, Max Welling, Roohollah Amiri, Srinivas Yerramalli, and Taesang Yoo. Neural rf slam for unsupervised positioning and mapping with channel state information, 2022.
- Ilia Karmanov, Farhad G Zanjani, Ishaque Kadampot, Simone Merlin, and Daniel Dijkman. Wicluster: Passive indoor 2d/3d positioning using wifi without precise labels. In *2021 IEEE Global Communications Conference (GLOBECOM)*, pp. 1–7. IEEE, 2021.
- Zahra Ezzati Khatab, Amirhosein Hajihoseini, and Seyed Ali Ghorashi. A fingerprint method for indoor localization using autoencoder based deep extreme learning machine. *IEEE sensors letters*, 2(1):1–4, 2017.
- Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Bal-subramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, pp. 5637–5664. PMLR, 2021.
- Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. *Advances in neural information processing systems*, 31, 2018.
- Jing Liu, Nan Liu, Zhiwen Pan, and Xiaohu You. Autloc: Deep autoencoder for indoor localization with rss fingerprinting. In *2018 10th International Conference on Wireless Communications and Signal Processing (WCSP)*, pp. 1–6, 2018. doi: 10.1109/WCSP.2018.8555665.
- Neal Patwari, Joshua N Ash, Spyros Kyperountas, Alfred O Hero, Randolph L Moses, and Neiyer S Correal. Locating the nodes: cooperative localization in wireless sensor networks. *IEEE Signal processing magazine*, 22(4):54–69, 2005.
- Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport. *Center for Research in Economics and Statistics Working Papers*, (2017-86), 2017.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pp. 8748–8763. PMLR, 2021.
- James Scott and Mike Hazas. User-friendly surveying techniques for location-aware systems. In *International Conference on Ubiquitous Computing*, pp. 44–53. Springer, 2003.

- Christoph Studer, Saïd Medjkouh, Emre Gonultas, Tom Goldstein, and Olav Tirkkonen. Channel charting: Locating users within the radio environment using channel state information. *IEEE Access*, 6:47682–47698, 2018.
- David Tse and Pramod Viswanath. *Fundamentals of wireless communication*. Cambridge university press, 2005.
- Olivia Wiles, Sven Gowal, Florian Stimberg, Sylvestre Alvisé-Rebuffi, Ira Ktena, Taylan Cemgil, et al. A fine-grained analysis on distribution shift. *arXiv preprint arXiv:2110.11328*, 2021.
- Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7959–7971, 2022.
- Farhad G Zanjani, Ilia Karmanov, Hanno Ackermann, Daniel Dijkman, Simone Merlin, Ishaque Kadampot, Brian Buesker, Vamsi Vegunta, and Fatih Porikli. Deep learning frameworks for weakly-supervised indoor localization. In *NeurIPS 2021 Competitions and Demonstrations Track*, pp. 349–354. PMLR, 2022.
- Mingmin Zhao, Tianhong Li, Mohammad Abu Alsheikh, Yonglong Tian, Hang Zhao, Antonio Torralba, and Dina Katabi. Through-wall human pose estimation using radio signals. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7356–7365, 2018a.
- Mingmin Zhao, Yonglong Tian, Hang Zhao, Mohammad Abu Alsheikh, Tianhong Li, Rumen Hristov, Zachary Kabelac, Dina Katabi, and Antonio Torralba. RF-based 3d skeletons. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, pp. 267–281, 2018b.

A BACKGROUND ON RF LOCALISATION

ToA and AoA estimates are traditionally arrived at using super-resolution methods (Henninger et al., 2022). maximum likelihood estimator (MLE) then uses ToA and AoA jointly in order to obtain position estimates.

Positioning model. We assume a positioning system comprised of M synchronised locators that listen for user devices. Each m th locator has known 3D position vector $\mathbf{p}_m = [x_m, y_m, z_m]$, and 3D 3×3 orientation matrix $\mathbf{\Omega}_m$. The estimation parameter is the user position $\mathbf{x} = [x, y, z]$. $\mathbf{\Omega}_m$ is formed by the orthogonal unit vectors corresponding to the m th locator’s orientation w.r.t. the reference coordinate system. Each k th locator can estimate the user’s distance \hat{d}_m that corresponds to its ToA and the impinging angle-of-arrival (AoA). ToA estimates implicitly include an unknown transmit time τ , since user devices are asynchronous w.r.t. the locators.

Joint ToA-AoA MLE. Assuming independent ToA and AoA estimates, which can be combined to obtain a more robust joint ToA-AoA position estimate

$$(\hat{\mathbf{x}}_\cap, \hat{\tau}_\cap) = \arg \max_{\mathbf{x}, \tau} \left\{ \sum_{m=1}^M w_{T,m} \cdot \ln \left\{ \frac{1}{\sqrt{2\pi}} \cdot \exp \left[- \frac{(\hat{d}_m - \|\mathbf{p}_m - \mathbf{x}\| - \tau \cdot c)^2}{2\sigma_m^2} \right] \right\} \right\} \quad (2)$$

$$+ \sum_m \kappa_m \hat{\mathbf{u}}_m^T \mathbf{\Omega}_m^T \frac{\mathbf{x} - \mathbf{p}_m}{\|\mathbf{x} - \mathbf{p}_m\|} \} \quad (3)$$

where σ_m^2 and $w_{T,m}$ are respectively the error variance and optimisation weight that correspond to the k th locator, τ denotes the unknown transmit time, and c the speed of light. where $\hat{\mathbf{u}}_m \in \mathbb{R}^3$ is a unit vector estimate of AoA in the locator’s reference frame $\mathbf{\Omega}_m$. The concentration parameters κ_m reflect the reliability of the angular measurements.

A.1 DEEP LEARNING

Due to the high dimensionality of radio frequency (RF) data deep learning (DL) techniques emerged to close the gap of localization systems in rich scattering environments.

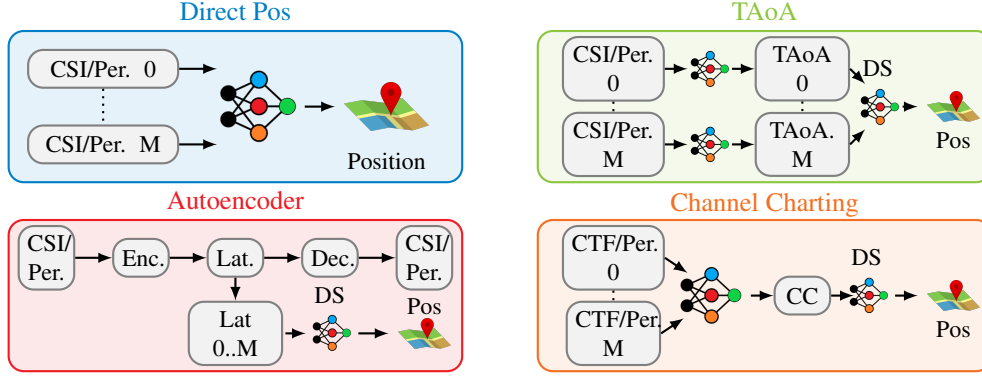


Figure 8: Different procedures to estimate the position from RF data.

Fig. 8 shows an extract of the currently implemented NN configurations, which are partly supervised and unsupervised with an additional delay spread (DS) task to predict the device position. The methods can be split into different types of prediction challenges. The first proposed techniques tried to estimate directly from the *raw* channel state information (CSI) the position directly, expecting the Neural Network (NN) to capture the channel model and invert it to the position. Another type is predicting the time-and-angle-of-arrival (TAoA) in a local coordinate system to remove any global structure and therefore a robust feature. An alternative to this low dimensional feature is the auto-encoder (AE) which allows to capture the spatial consistency of the channel in the latent space. The currently most favoured approach is channel-charting (CC) where a low-dimesional channel chart is estimated, by forcing a spatial consistency in the channel chart.

A.2 SUPERVISED LEARNING

Supervised learning was the first approach to be tackled, where a huge amount of costly labels were created using advanced Light Detection and Ranging (LiDaR) systems. We consider three different strains:

1.) Directly to postion from high dimensionality CSI/PER

Using this high dimensional information the direct mapping,

$$\mathcal{C} : \mathbb{C}^M \rightarrow \mathbb{R}^3$$

is learned, where fig. 8 part one shows this concept. It was shown that this approach generally results in a high dimensionality fingerprinting system.

2.) Predicting TAoA

Predicting the TAoA from the raw information allows to learn a similiar mapping

$$\mathcal{C} : \mathbb{C}^M \rightarrow \mathbb{R}^3$$

but converting the output dimensionality in a local coordinate system independent from other influences. To convert the TAoA information into a position a classical weighted system (Henninger et al., 2022) could be used or as shown in fig. 8 part two a downstream NN.

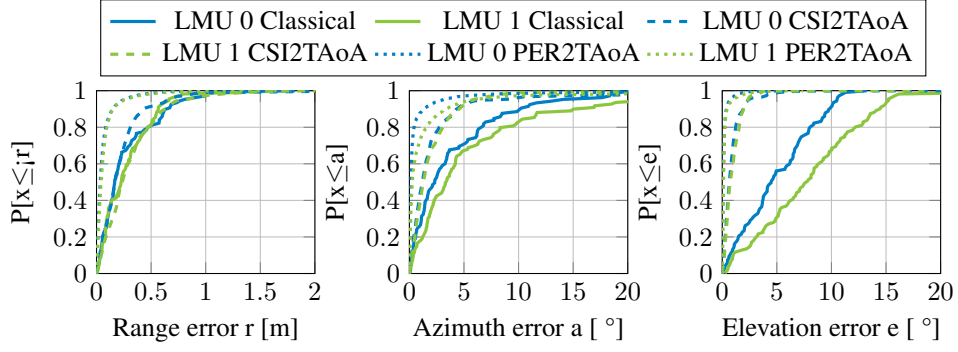


Figure 9: Example prediction of local coordinates using the classical or the proposed NNs.

Fig. 9 depicts that NN can beat due to resilience against hardware impairments (non ideal antenna aperture and oscillators) and rich scattering, where the classical approach crumbles.

3.) Directly to position from TAoA An alternative flavour is to predict from the TAoA the position. Using multiple locators results in a weighting based on the position system, and thus a form of direct memorization between regions.

A.3 UNSUPERVISED LEARNING

Due to the costs of creating precise labels as well as updating the labels if the environment changes, multiple unsupervised approaches hoping to generalize better were proposed. We highlight only two currently most favoured methods.

A.3.1 CHANNEL CHARTING

CC exploits the spatial consistency of the channel by first normalizing it with the path-loss via

$$\tilde{\mathbf{H}} = \frac{B^{\beta-1}}{\|\tilde{\mathbf{H}}\|_F^\beta} \tilde{\mathbf{H}}$$

where B is the number of receive antennas, β the path-loss exponent and $\|\cdot\|_F$ the forbenius norm, respectively. As this technique relies on distilling the spatial consistency by a (hopefully) loss-less dimensionality reduction (e.g. averaging the number of subcarriers), the channel chart can be learned by the mapping

$$\mathcal{C} : \mathbb{C}^{M'} \rightarrow \mathbb{R}^D,$$

where D is typically a low dimensional vector (normally two). This channel chart is learned using the triplet loss.

As this method is fully unsupervised an the mapping is not corresponding to the actual ground-truth three different metrics were introduced: (i) continuity (CT) measuring for the K neighbors in the original space the point-wise continuity, e.g. the points are following the same order in the original space as well as the latent space (ii) trustworthiness (TW) measuring the point-wise trustworthiness of the point surrounding in the latent space, e.g. the false points close to the original data point reduces this metric. For more details we refer to (Studer et al., 2018)

A.3.2 AUTOENCODER

The basic idea of an AE is to learn two functions, an encoder \mathcal{C} and a decoder \mathcal{C}^{-1} , so that the average approximation error

$$E = \frac{1}{N} \sum_{n=1}^N \|\mathbf{f}_n - \mathcal{C}^{-1}(\mathcal{C}(\mathbf{f}_n))\|_2^2 \quad (4)$$

for a set of vectors \mathbf{f}_n $n = 1$ is minimal. The hope is that the AE creates a low dimensional representation, capturing the essential components of the channel. These components are later exploited for a downstream task.

Table 4: TAoA2Pos

| Parameter | Type | Choices | Final |
|---------------|---------|--|---------------------|
| epochs | choice | [20, 30, 40] | 20 |
| lr | choice | [0.0001, 0.001, 0.01] | 0.0001 |
| momentum | uniform | [0.8, 1] | 0.8392178101801823 |
| step size | choice | [1, 2, 10] | 10 |
| gamma | uniform | [0.4, 0.8] | 0.41178097722520307 |
| act func | choice | ['ReLU', 'LeakyReLU', 'Sigmoid', 'Tanh', 'Softplus'] | ReLU |
| last act func | choice | ['Sigmoid'] | Sigmoid |
| optimizer | choice | ['SGD', 'ADAM'] | SGD |
| loss func | choice | ['MSE', 'L1'] | MSE |
| lin1 size | choice | [32, 64, 128, 256, 512] | 512 |
| lin2 size | choice | [32, 64, 128, 256, 512] | 128 |
| lin3 size | choice | [32, 64, 128, 256, 512] | 128 |

Table 5: CSI2Pos

| Parameter | Type | Choices | Final |
|---------------|---------|--|---------------------|
| epochs | choice | [20, 30, 40] | 20 |
| lr | choice | [0.0001, 0.001, 0.01] | 0.0001 |
| momentum | uniform | [0.8, 1] | 0.8853297100815697 |
| step size | choice | [1, 2, 10] | 10 |
| gamma | uniform | [0.4, 0.8] | 0.43883862858760825 |
| act func | choice | ['ReLU', 'LeakyReLU', 'Sigmoid', 'Tanh', 'Softplus'] | ReLU |
| last act func | choice | ['Sigmoid'] | Sigmoid |
| optimizer | choice | ['SGD', 'ADAM'] | SGD |
| loss func | choice | ['MSE', 'L1'] | L1 |
| branch1 | choice | [0, 1] | 1 |
| branch2 | choice | [0, 1] | 1 |
| c1 size | choice | [8, 16, 32, 64] | 8 |
| c2 size | choice | [8, 16, 32, 64] | 16 |
| c3 size | choice | [8, 16, 32, 64] | 64 |
| c4 size | choice | [8, 16, 32, 64] | 8 |
| c5 size | choice | [8, 16, 32, 64] | 64 |
| c6 size | choice | [8, 16, 32, 64] | 32 |
| c7 size | choice | [8, 16, 32, 64] | 16 |
| c8 size | choice | [8, 16, 32, 64] | 8 |
| c9 size | choice | [8, 16, 32, 64] | 16 |
| c10 size | choice | [8, 16, 32, 64] | 16 |
| k2 size | choice | [1, 2, 3, 4, 5, 6] | 4 |
| k4 size | choice | [1, 2, 3, 4, 5, 6] | 3 |
| k6 size | choice | [1, 2, 3, 4, 5, 6] | 4 |
| k8 size | choice | [1, 2, 3, 4, 5, 6] | 3 |
| k10 size | choice | [1, 2, 3, 4, 5, 6] | 5 |
| k12 size | choice | [1, 2, 3, 4, 5, 6] | 2 |
| k14 size | choice | [1, 2, 3, 4, 5, 6] | 3 |
| k16 size | choice | [1, 2, 3, 4, 5, 6] | 6 |
| k18 size | choice | [1, 2, 3, 4, 5, 6] | 6 |
| k20 size | choice | [1, 2, 3, 4, 5, 6] | 3 |
| s2 size | choice | [1, 2, 3, 4, 5] | 2 |
| s4 size | choice | [1, 2, 3, 4, 5, 6] | 1 |
| s6 size | choice | [1, 2, 3, 4, 5, 6] | 1 |
| s8 size | choice | [1, 2, 3, 4, 5, 6] | 3 |
| s10 size | choice | [1, 2, 3, 4, 5, 6] | 2 |
| s12 size | choice | [1, 2, 3, 4, 5] | 4 |
| s14 size | choice | [1, 2, 3, 4, 5, 6] | 1 |
| s16 size | choice | [1, 2, 3, 4, 5, 6] | 5 |
| s18 size | choice | [1, 2, 3, 4, 5, 6] | 5 |
| s20 size | choice | [1, 2, 3, 4, 5, 6] | 1 |
| m2 size | choice | [1, 2, 3, 4, 5, 6] | 6 |
| m4 size | choice | [1, 2, 3, 4, 5, 6] | 3 |
| m6 size | choice | [1, 2, 3, 4, 5, 6] | 1 |
| m8 size | choice | [1, 2, 3, 4, 5, 6] | 2 |
| m10 size | choice | [1, 2, 3, 4, 5, 6] | 5 |
| m12 size | choice | [1, 2, 3, 4, 5] | 1 |
| m14 size | choice | [1, 2, 3, 4, 5, 6] | 6 |
| m16 size | choice | [1, 2, 3, 4, 5, 6] | 4 |
| m18 size | choice | [1, 2, 3, 4, 5, 6] | 5 |
| m20 size | choice | [1, 2, 3, 4, 5, 6] | 4 |

B DL MODELS

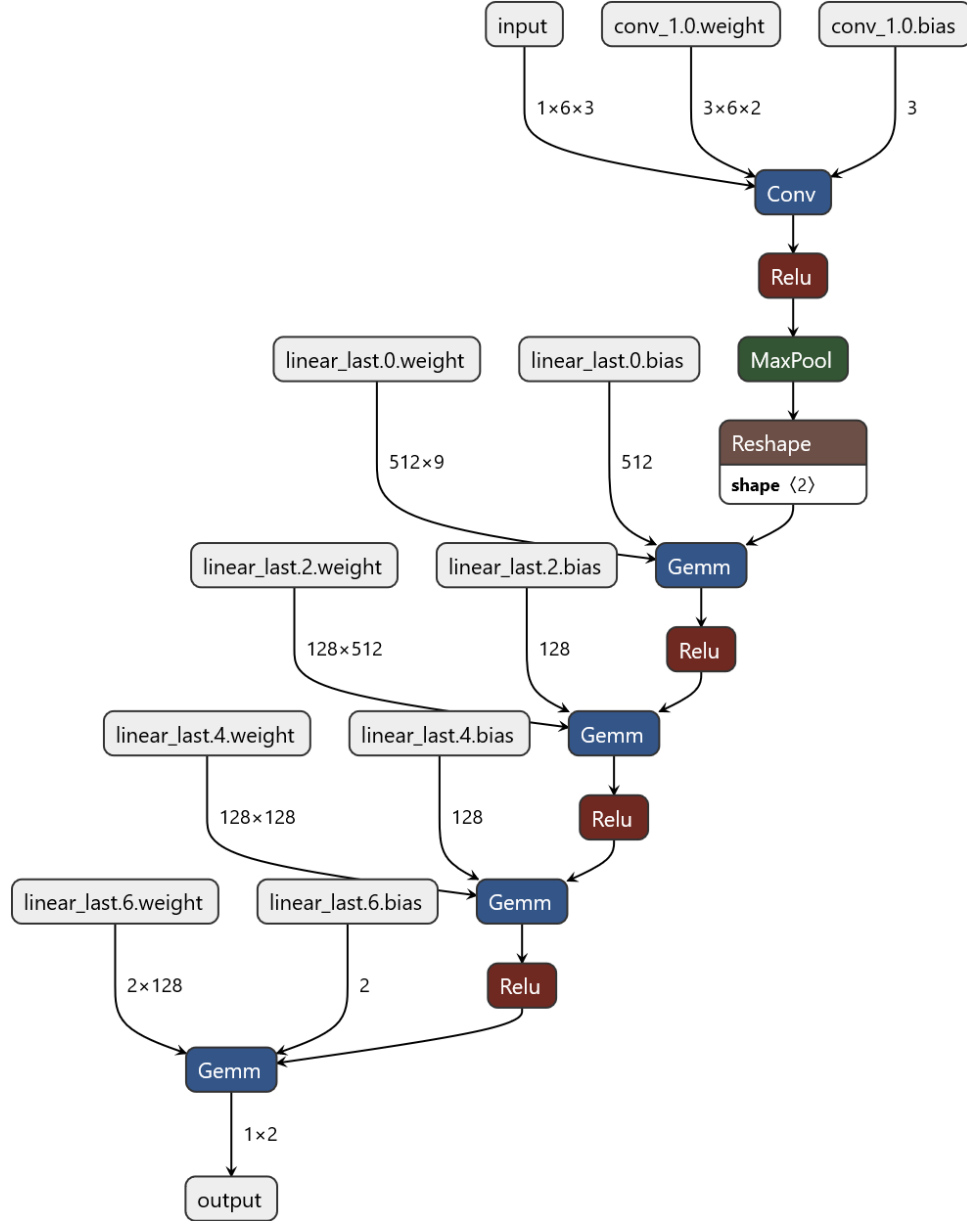


Figure 10: Structure TAoA2Pos

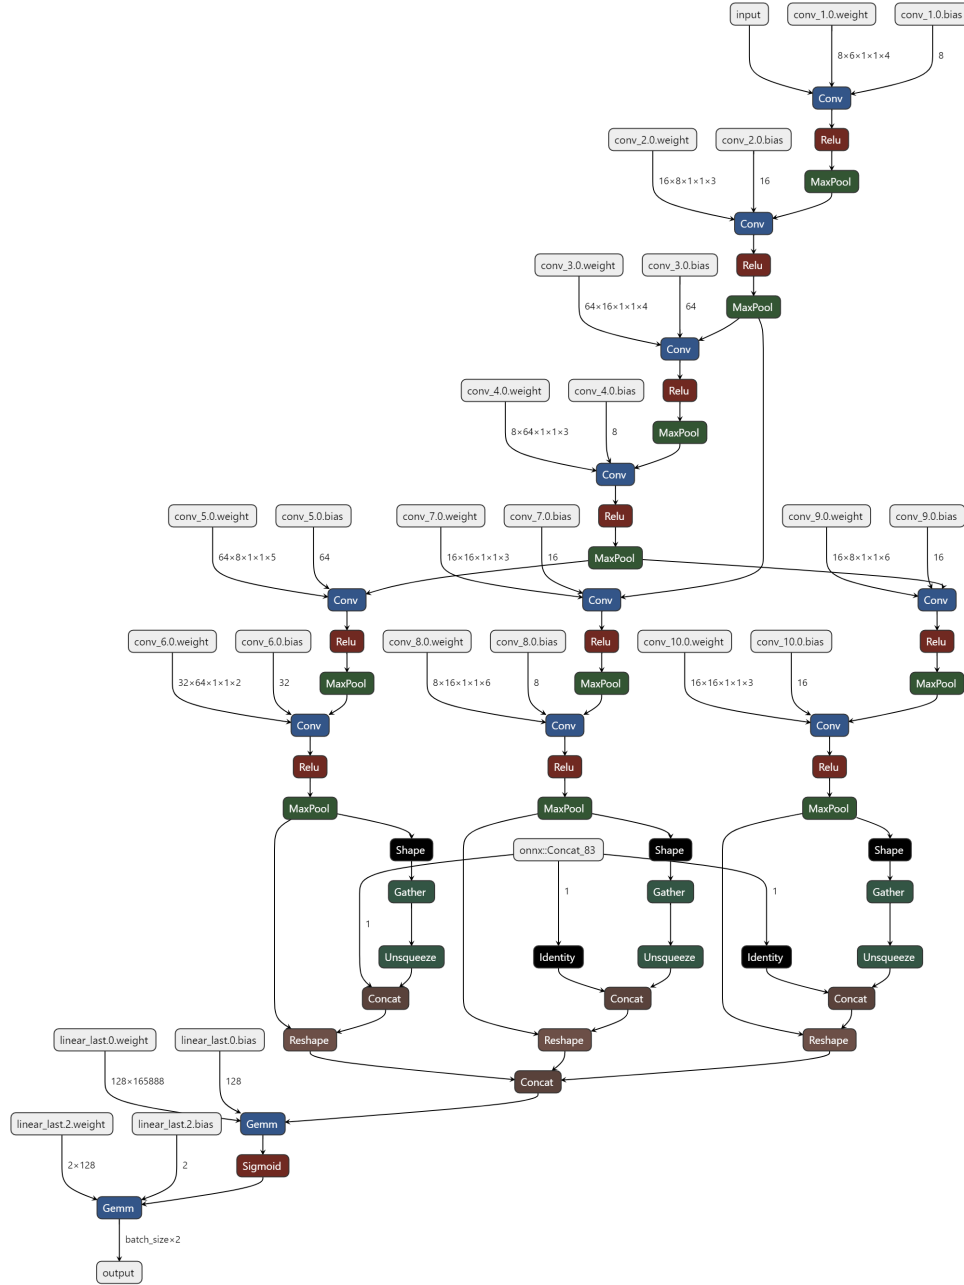


Figure 11: Structure CSI2Pos

Table 6: CSI2TAoA

| Parameter | Type | Choices | Final |
|---------------|---------|--|--------------------|
| epochs | choice | [20, 30, 40] | 30 |
| lr | choice | [0.0001, 0.001, 0.01] | 0.0001 |
| momentum | uniform | [0.8, 1] | 0.960095906476885 |
| step size | choice | [1, 2, 10] | 10 |
| gamma | uniform | [0.4, 0.8] | 0.7152406448940722 |
| act func | choice | ['ReLU', 'LeakyReLU', 'Sigmoid', 'Tanh', 'Softplus'] | Tanh |
| last act func | choice | ['Sigmoid'] | Sigmoid |
| optimizer | choice | ['SGD', 'ADAM'] | ADAM |
| loss func | choice | ['MSE', 'L1'] | L1 |
| branch1 | choice | [0, 1] | 1 |
| branch2 | choice | [0, 1] | 0 |
| c1 size | choice | [8, 16, 32, 64] | 8 |
| c2 size | choice | [8, 16, 32, 64] | 32 |
| c3 size | choice | [8, 16, 32, 64] | 8 |
| c4 size | choice | [8, 16, 32, 64] | 16 |
| c5 size | choice | [8, 16, 32, 64] | 8 |
| c6 size | choice | [8, 16, 32, 64] | 64 |
| c7 size | choice | [8, 16, 32, 64] | 16 |
| c8 size | choice | [8, 16, 32, 64] | 32 |
| c9 size | choice | [8, 16, 32, 64] | 64 |
| c10 size | choice | [8, 16, 32, 64] | 16 |
| k1 size | choice | [1] | 1 |
| k2 size | choice | [1, 2, 3, 4, 5, 6] | 1 |
| k3 size | choice | [1] | 1 |
| k4 size | choice | [1, 2, 3, 4, 5, 6] | 1 |
| k5 size | choice | [1] | 1 |
| k6 size | choice | [1, 2, 3, 4, 5, 6] | 3 |
| k7 size | choice | [1] | 1 |
| k8 size | choice | [1, 2, 3, 4, 5, 6] | 6 |
| k9 size | choice | [1] | 1 |
| k10 size | choice | [1, 2, 3, 4, 5, 6] | 2 |
| k11 size | choice | [1] | 1 |
| k12 size | choice | [1, 2, 3, 4, 5, 6] | 2 |
| k13 size | choice | [1] | 1 |
| k14 size | choice | [1, 2, 3, 4, 5, 6] | 3 |
| k15 size | choice | [1] | 1 |
| k16 size | choice | [1, 2, 3, 4, 5, 6] | 3 |
| k17 size | choice | [1] | 1 |
| k18 size | choice | [1, 2, 3, 4, 5, 6] | 5 |
| k19 size | choice | [1] | 1 |
| k20 size | choice | [1, 2, 3, 4, 5, 6] | 2 |
| s1 size | choice | [1] | 1 |
| s2 size | choice | [1, 2, 3, 4, 5] | 1 |
| s3 size | choice | [1] | 1 |
| s4 size | choice | [1, 2, 3, 4, 5, 6] | 2 |
| s5 size | choice | [1] | 1 |
| s6 size | choice | [1, 2, 3, 4, 5, 6] | 4 |
| s7 size | choice | [1] | 1 |
| s8 size | choice | [1, 2, 3, 4, 5, 6] | 6 |
| s9 size | choice | [1] | 1 |
| s10 size | choice | [1, 2, 3, 4, 5, 6] | 3 |
| s11 size | choice | [1] | 1 |
| s12 size | choice | [1, 2, 3, 4, 5] | 1 |
| s13 size | choice | [1] | 1 |
| s14 size | choice | [1, 2, 3, 4, 5, 6] | 3 |
| s15 size | choice | [1] | 1 |
| s16 size | choice | [1, 2, 3, 4, 5, 6] | 6 |
| s17 size | choice | [1] | 1 |
| s18 size | choice | [1, 2, 3, 4, 5, 6] | 3 |
| s19 size | choice | [1] | 1 |
| s20 size | choice | [1, 2, 3, 4, 5, 6] | 3 |
| m1 size | choice | [1] | 1 |
| m2 size | choice | [1, 2, 3, 4, 5, 6] | 4 |
| m3 size | choice | [1] | 1 |
| m4 size | choice | [1, 2, 3, 4, 5, 6] | 4 |
| m5 size | choice | [1] | 1 |
| m6 size | choice | [1, 2, 3, 4, 5, 6] | 2 |
| m7 size | choice | [1] | 1 |
| m8 size | choice | [1, 2, 3, 4, 5, 6] | 1 |
| m9 size | choice | [1] | 1 |
| m10 size | choice | [1, 2, 3, 4, 5, 6] | 4 |
| m11 size | choice | [1] | 1 |
| m12 size | choice | [1, 2, 3, 4, 5] | 3 |
| m13 size | choice | [1] | 1 |
| m14 size | choice | [1, 2, 3, 4, 5, 6] | 2 |
| m15 size | choice | [1] | 1 |
| m16 size | choice | [1, 2, 3, 4, 5, 6] | 6 |
| m17 size | choice | [1] | 1 |
| m18 size | choice | [1, 2, 3, 4, 5, 6] | 5 |
| m19 size | choice | [1] | 1 |
| m20 size | choice | [1, 2, 3, 4, 5, 6] | 6 |

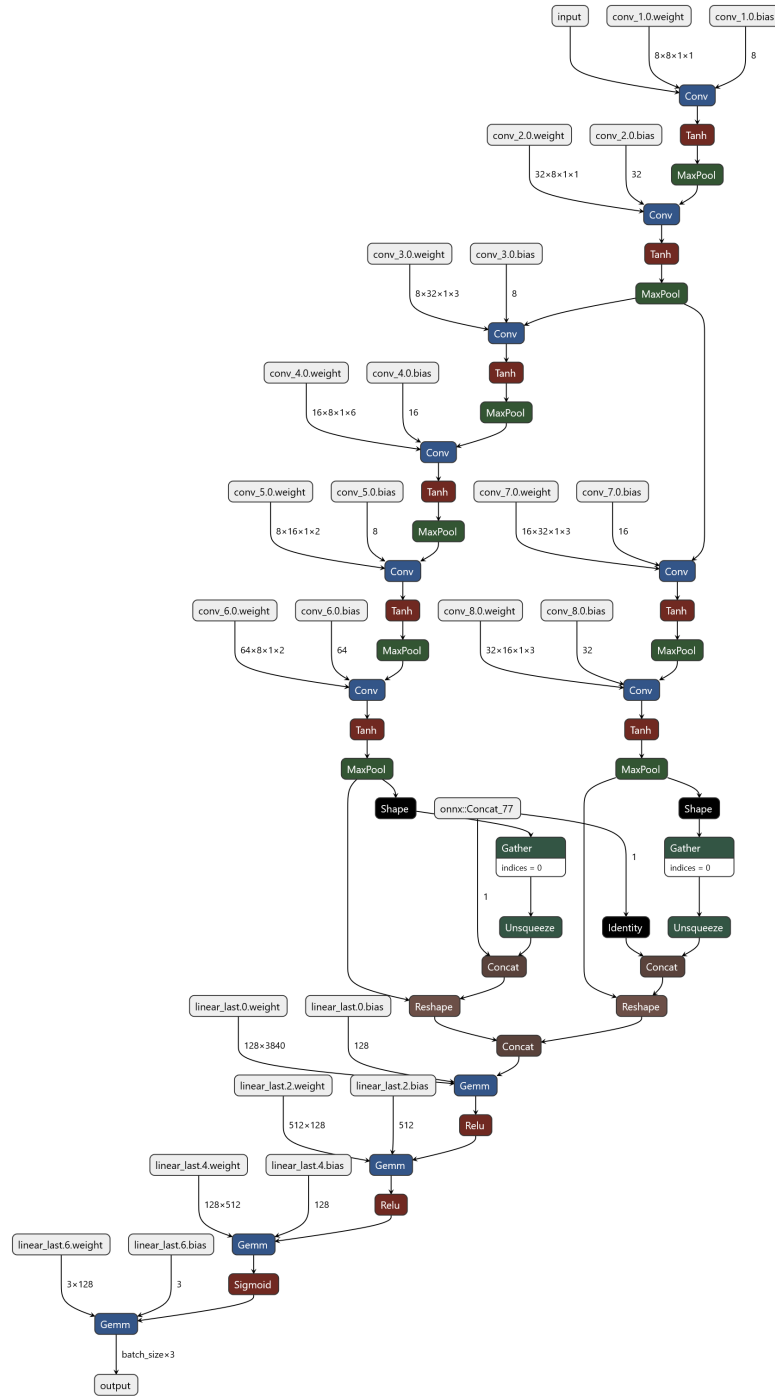


Figure 12: Structure CSI2TAoA

Table 7: CSI AE

| Parameter | Type | Choices | Final |
|---------------|---------|--|---------------------|
| epochs | choice | [20, 30, 40] | 20 |
| lr | choice | [0.0001, 0.001, 0.01] | 0.0001 |
| momentum | uniform | [0.8, 1] | 0.9363040304887302 |
| step size | choice | [1, 2, 10] | 2 |
| gamma | uniform | [0.4, 0.8] | 0.44213564205749145 |
| act func | choice | ['ReLU', 'LeakyReLU', 'Sigmoid', 'Tanh', 'Softplus'] | ReLU |
| last act func | choice | ['Sigmoid'] | Sigmoid |
| optimizer | choice | ['SGD', 'ADAM'] | ADAM |
| loss func | choice | ['MSE', 'L1', 'NMSE', 'BCE'] | L1 |
| lin1 size | choice | [32, 64, 128, 256, 512] | 256 |
| lin2 size | choice | [32, 64, 128, 256, 512] | 64 |
| lin3 size | choice | [32, 64, 128, 256, 512] | 128 |
| lin4 size | choice | [32, 64, 128, 256, 512] | 128 |
| feat dim | choice | [128, 256, 512, 1024] | 256 |

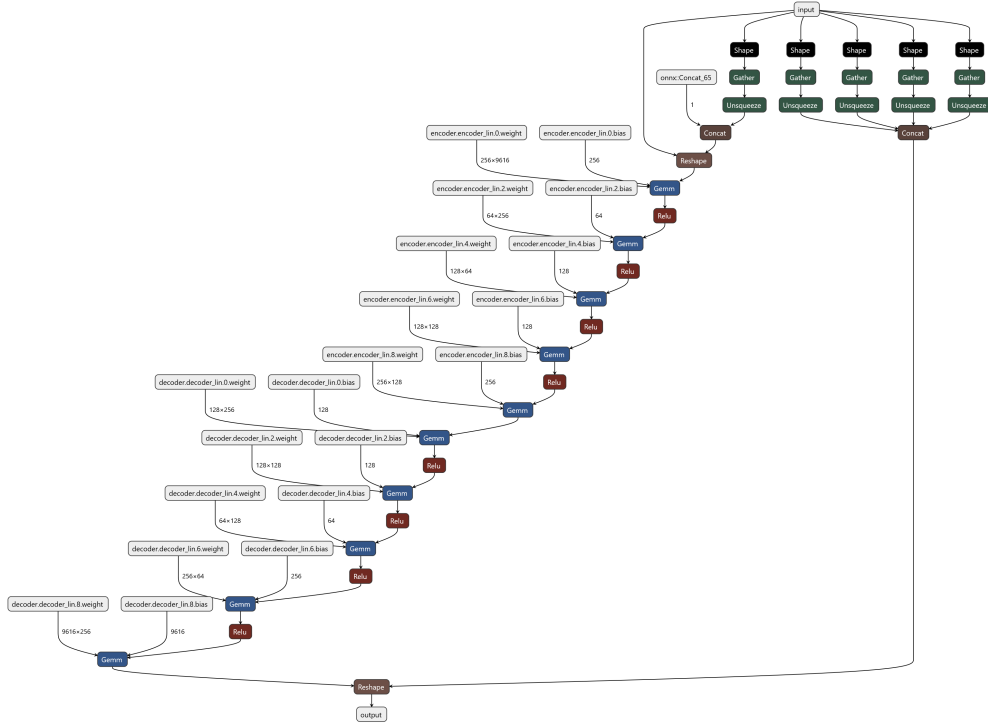


Figure 13: Structure CSI AE

Table 8: PER2Pos

| Parameter | Type | Choices | Final |
|---------------|---------|--|--------------------|
| epochs | choice | [20, 30, 40] | 40 |
| lr | choice | [0.0001, 0.001, 0.01] | 0.001 |
| momentum | uniform | [0.8, 1] | 0.9997778126234411 |
| step size | choice | [1, 2, 10] | 1 |
| gamma | uniform | [0.4, 0.8] | 0.560402530889289 |
| act func | choice | ['ReLU', 'LeakyReLU', 'Sigmoid', 'Tanh', 'Softplus'] | ReLU |
| last act func | choice | ['Sigmoid'] | Sigmoid |
| optimizer | choice | ['SGD', 'ADAM'] | ADAM |
| loss func | choice | ['MSE', 'L1'] | MSE |
| branch1 | choice | [0, 1] | 1 |
| branch2 | choice | [0, 1] | 0 |
| c1 size | choice | [8, 16, 32, 64] | 64 |
| c2 size | choice | [8, 16, 32, 64] | 32 |
| c3 size | choice | [8, 16, 32, 64] | 16 |
| c4 size | choice | [8, 16, 32, 64] | 8 |
| c5 size | choice | [8, 16, 32, 64] | 32 |
| c6 size | choice | [8, 16, 32, 64] | 64 |
| c7 size | choice | [8, 16, 32, 64] | 8 |
| c8 size | choice | [8, 16, 32, 64] | 32 |
| c9 size | choice | [8, 16, 32, 64] | 8 |
| c10 size | choice | [8, 16, 32, 64] | 8 |
| k1 size | choice | [1, 2, 4] | 1 |
| k2 size | choice | [1, 2, 4] | 1 |
| k3 size | choice | [1, 2] | 1 |
| k4 size | choice | [1, 2, 4] | 4 |
| k5 size | choice | [1, 2] | 2 |
| k6 size | choice | [1, 2, 4] | 4 |
| k7 size | choice | [1, 2] | 2 |
| k8 size | choice | [1, 2, 4] | 1 |
| k9 size | choice | [1, 2] | 1 |
| k10 size | choice | [1, 2, 4] | 1 |
| k11 size | choice | [1, 2] | 2 |
| k12 size | choice | [1, 2, 4] | 4 |
| k13 size | choice | [1, 2] | 2 |
| k14 size | choice | [1, 2, 4] | 2 |
| k15 size | choice | [1, 2] | 2 |
| k16 size | choice | [1, 2, 4] | 1 |
| k17 size | choice | [1, 2] | 1 |
| k18 size | choice | [1, 2, 4] | 2 |
| k19 size | choice | [1, 2] | 2 |
| k20 size | choice | [1, 2, 4] | 4 |
| s1 size | choice | [1, 2] | 4 |
| s2 size | choice | [1, 2, 3, 4, 5] | 2 |
| s3 size | choice | [1, 2] | 4 |
| s4 size | choice | [1, 2, 4] | 1 |
| s5 size | choice | [1, 2] | 4 |
| s6 size | choice | [1, 2, 4] | 2 |
| s7 size | choice | [1, 2] | 4 |
| s8 size | choice | [1, 2, 4] | 4 |
| s9 size | choice | [1, 2] | 1 |
| s10 size | choice | [1, 2, 4] | 4 |
| s11 size | choice | [1, 2] | 1 |
| s12 size | choice | [1, 2, 3, 4, 5] | 1 |
| s13 size | choice | [1, 2] | 2 |
| s14 size | choice | [1, 2, 4] | 4 |
| s15 size | choice | [1, 2] | 1 |
| s16 size | choice | [1, 2, 4] | 1 |
| s17 size | choice | [1, 2] | 2 |
| s18 size | choice | [1, 2, 4] | 4 |
| s19 size | choice | [1, 2] | 8 |
| s20 size | choice | [1, 2, 4] | 8 |
| m1 size | choice | [1, 2] | 2 |
| m2 size | choice | [1, 2, 4] | 2 |
| m3 size | choice | [1, 2] | 1 |
| m4 size | choice | [1, 2, 4] | 4 |
| m5 size | choice | [1, 2] | 1 |
| m6 size | choice | [1, 2, 4] | 2 |
| m7 size | choice | [1, 2] | 1 |
| m8 size | choice | [1, 2, 4] | 1 |
| m9 size | choice | [1, 2] | 1 |
| m10 size | choice | [1, 2, 4] | 2 |
| m11 size | choice | [1, 2] | 1 |
| m12 size | choice | [1, 2, 3, 4, 5] | 2 |
| m13 size | choice | [1, 2] | 1 |
| m14 size | choice | [1, 2, 4] | 2 |
| m15 size | choice | [1, 2] | 1 |
| m16 size | choice | [1, 2, 4] | 2 |
| m17 size | choice | [1, 2] | 1 |
| m18 size | choice | [1, 2, 4] | 4 |
| m19 size | choice | [1, 2] | 2 |
| m20 size | choice | [1, 2, 4] | 4 |

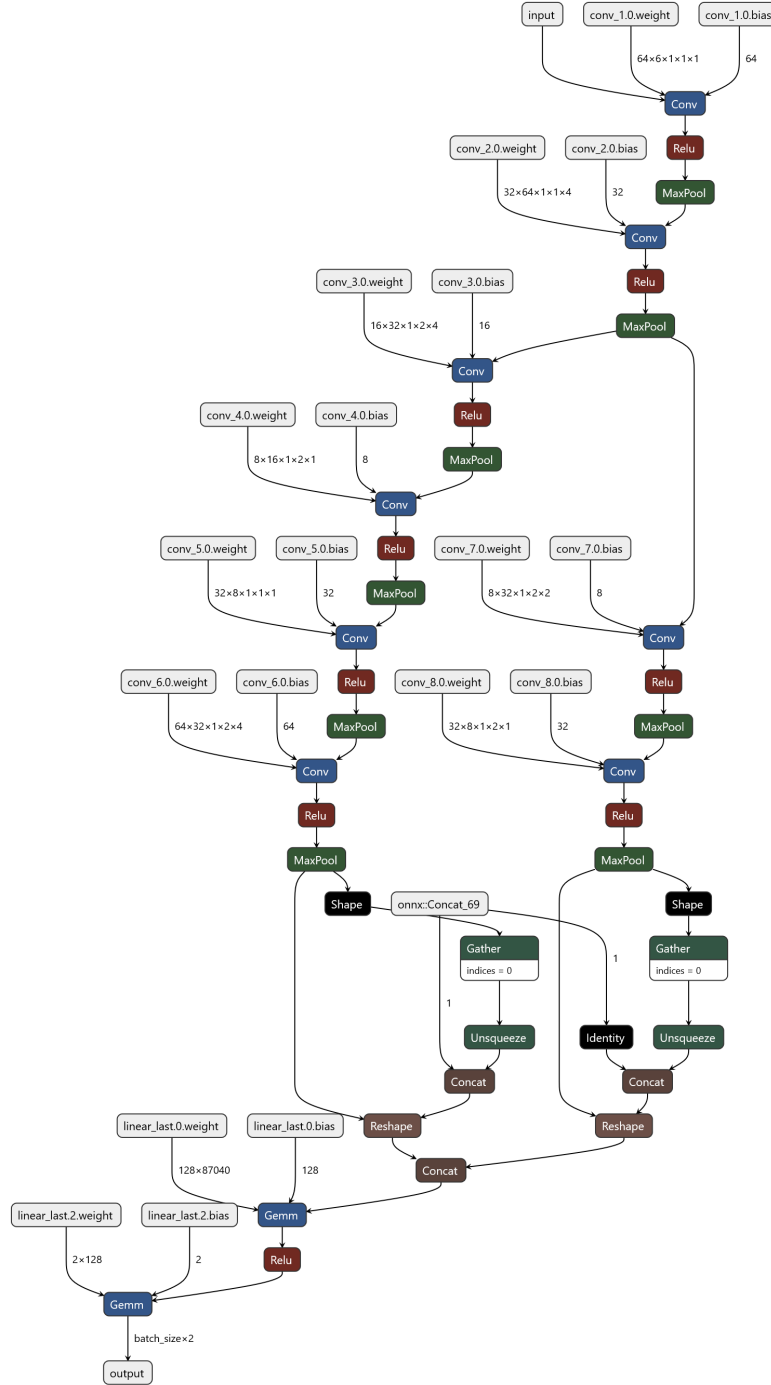


Figure 14: Structure Per2Pos

Table 9: PER2TAoA

| Parameter | Type | Choices | Final |
|---------------|---------|--|--------------------|
| epochs | choice | [20, 30, 40] | 30 |
| lr | choice | [0.0001, 0.001, 0.01] | 0.0001 |
| momentum | uniform | [0.8, 1] | 0.8999771834863867 |
| step size | choice | [1, 2, 10] | 2 |
| gamma | uniform | [0.4, 0.8] | 0.7350887010069965 |
| act func | choice | ['ReLU', 'LeakyReLU', 'Sigmoid', 'Tanh', 'Softplus'] | Softplus |
| last act func | choice | ['Sigmoid'] | Sigmoid |
| optimizer | choice | ['SGD', 'ADAM'] | SGD |
| loss func | choice | ['MSE', 'L1'] | L1 |
| branch1 | choice | [0, 1] | 0 |
| branch2 | choice | [0, 1] | 0 |
| c1 size | choice | [8, 16, 32, 64] | 16 |
| c2 size | choice | [8, 16, 32, 64] | 16 |
| c3 size | choice | [8, 16, 32, 64] | 8 |
| c4 size | choice | [8, 16, 32, 64] | 16 |
| c5 size | choice | [8, 16, 32, 64] | 64 |
| c6 size | choice | [8, 16, 32, 64] | 16 |
| c7 size | choice | [8, 16, 32, 64] | 16 |
| c8 size | choice | [8, 16, 32, 64] | 16 |
| c9 size | choice | [8, 16, 32, 64] | 16 |
| c10 size | choice | [8, 16, 32, 64] | 8 |
| k1 size | choice | [1, 2, 4] | 2 |
| k2 size | choice | [1, 2, 4] | 1 |
| k3 size | choice | [1, 2] | 1 |
| k4 size | choice | [1, 2, 4] | 2 |
| k5 size | choice | [1, 2] | 1 |
| k6 size | choice | [1, 2, 4] | 1 |
| k7 size | choice | [1, 2] | 2 |
| k8 size | choice | [1, 2, 4] | 1 |
| k9 size | choice | [1, 2] | 2 |
| k10 size | choice | [1, 2, 4] | 1 |
| k11 size | choice | [1, 2] | 1 |
| k12 size | choice | [1, 2, 4] | 2 |
| k13 size | choice | [1, 2] | 2 |
| k14 size | choice | [1, 2, 4] | 4 |
| k15 size | choice | [1, 2] | 1 |
| k16 size | choice | [1, 2, 4] | 4 |
| k17 size | choice | [1, 2] | 2 |
| k18 size | choice | [1, 2, 4] | 2 |
| k19 size | choice | [1, 2] | 1 |
| k20 size | choice | [1, 2, 4] | 1 |
| s1 size | choice | [1, 2] | 2 |
| s2 size | choice | [1, 2, 3, 4, 5] | 2 |
| s3 size | choice | [1, 2] | 1 |
| s4 size | choice | [1, 2, 4] | 1 |
| s5 size | choice | [1, 2] | 1 |
| s6 size | choice | [1, 2, 4] | 4 |
| s7 size | choice | [1, 2] | 2 |
| s8 size | choice | [1, 2, 4] | 4 |
| s9 size | choice | [1, 2] | 2 |
| s10 size | choice | [1, 2, 4] | 1 |
| s11 size | choice | [1, 2] | 1 |
| s12 size | choice | [1, 2, 3, 4, 5] | 3 |
| s13 size | choice | [1, 2] | 2 |
| s14 size | choice | [1, 2, 4] | 4 |
| s15 size | choice | [1, 2] | 1 |
| s16 size | choice | [1, 2, 4] | 1 |
| s17 size | choice | [1, 2] | 1 |
| s18 size | choice | [1, 2, 4] | 1 |
| s19 size | choice | [1, 2] | 4 |
| s20 size | choice | [1, 2, 4] | 2 |
| m1 size | choice | [1, 2] | 1 |
| m2 size | choice | [1, 2, 4] | 4 |
| m3 size | choice | [1, 2] | 1 |
| m4 size | choice | [1, 2, 4] | 4 |
| m5 size | choice | [1, 2] | 2 |
| m6 size | choice | [1, 2, 4] | 4 |
| m7 size | choice | [1, 2] | 1 |
| m8 size | choice | [1, 2, 4] | 4 |
| m9 size | choice | [1, 2] | 1 |
| m10 size | choice | [1, 2, 4] | 2 |
| m11 size | choice | [1, 2] | 2 |
| m12 size | choice | [1, 2, 3, 4, 5] | 5 |
| m13 size | choice | [1, 2] | 1 |
| m14 size | choice | [1, 2, 4] | 2 |
| m15 size | choice | [1, 2] | 2 |
| m16 size | choice | [1, 2, 4] | 2 |
| m17 size | choice | [1, 2] | 2 |
| m18 size | choice | [1, 2, 4] | 2 |
| m19 size | choice | [1, 2] | 2 |
| m20 size | choice | [1, 2, 4] | 2 |

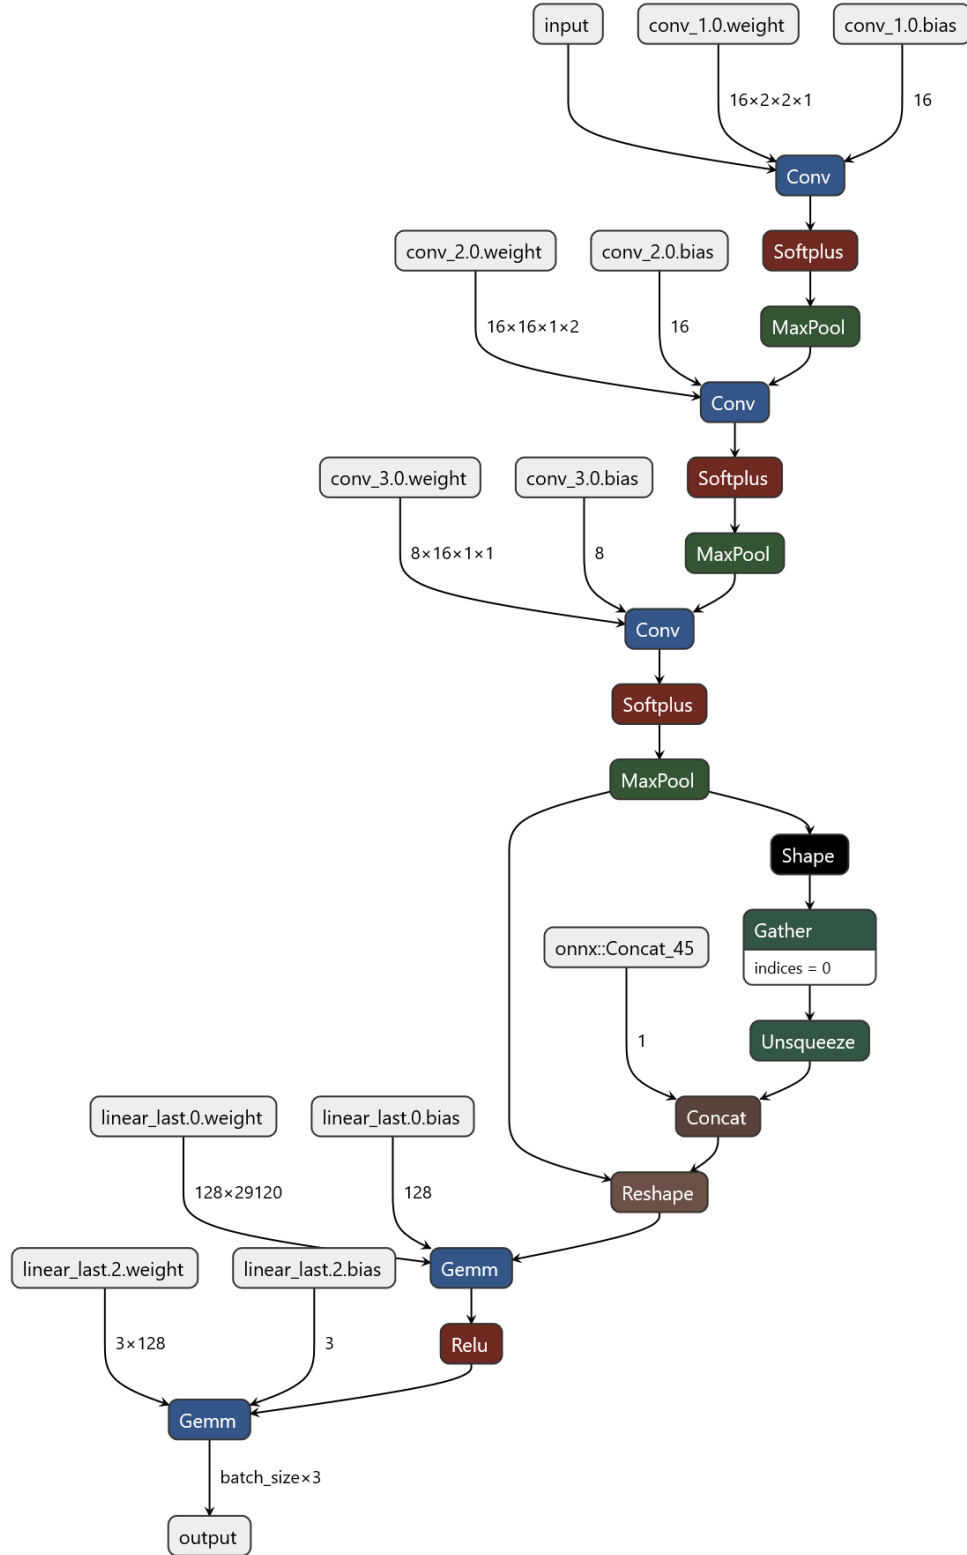


Figure 15: Structure Per2TAoA

Table 10: PER AE

| Parameter | Type | Choices | Final |
|---------------|---------|--|--------------------|
| epochs | choice | [20, 30, 40] | 30 |
| lr | choice | [0.0001, 0.001, 0.01] | 0.001 |
| momentum | uniform | [0.8, 1] | 0.9593149656519064 |
| step size | choice | [1, 2, 10] | 10 |
| gamma | uniform | [0.4, 0.8] | 0.870379744744098 |
| act func | choice | ['ReLU', 'LeakyReLU', 'Sigmoid', 'Tanh', 'Softplus'] | ReLU |
| last act func | choice | ['Sigmoid'] | Sigmoid |
| optimizer | choice | ['SGD', 'ADAM'] | ADAM |
| loss func | choice | ['MSE', 'L1', 'NMSE', 'BCE'] | BCE |
| branch1 | choice | [0, 1] | 1 |
| branch2 | choice | [0, 1] | 1 |
| c1 size | choice | [8, 16, 32, 64] | 64 |
| c2 size | choice | [8, 16, 32, 64] | 32 |
| c3 size | choice | [8, 16, 32, 64] | 16 |
| c4 size | choice | [1] | 8 |
| c5 size | choice | [1] | 32 |
| c6 size | choice | [1] | 64 |
| c7 size | choice | [1] | 8 |
| c8 size | choice | [1] | 32 |
| c9 size | choice | [1] | 8 |
| c10 size | choice | [1] | 64 |
| k1 size | choice | [1] | 1 |
| k2 size | choice | [1, 2, 4] | 1 |
| k3 size | choice | [1, 2] | 1 |
| k4 size | choice | [1, 2, 4] | 4 |
| k5 size | choice | [1, 2] | 2 |
| k6 size | choice | [1, 2, 4] | 4 |
| k7 size | choice | [1] | 2 |
| k8 size | choice | [1] | 1 |
| k9 size | choice | [1] | 1 |
| k10 size | choice | [1] | 1 |
| k11 size | choice | [1] | 2 |
| k12 size | choice | [1] | 4 |
| k13 size | choice | [1] | 2 |
| k14 size | choice | [1] | 2 |
| k15 size | choice | [1] | 2 |
| k16 size | choice | [1] | 1 |
| k17 size | choice | [1] | 1 |
| k18 size | choice | [1] | 2 |
| k19 size | choice | [1] | 2 |
| k20 size | choice | [1] | 4 |
| s1 size | choice | [1] | 2 |
| s2 size | choice | [1, 2, 4] | 2 |
| s3 size | choice | [1, 2] | 2 |
| s4 size | choice | [1, 2, 4] | 1 |
| s5 size | choice | [1, 2] | 1 |
| s6 size | choice | [1, 2, 4] | 2 |
| s7 size | choice | [1] | 2 |
| s8 size | choice | [1] | 2 |
| s9 size | choice | [1] | 1 |
| s10 size | choice | [1] | 4 |
| s11 size | choice | [1] | 1 |
| s12 size | choice | [1] | 1 |
| s13 size | choice | [1] | 2 |
| s14 size | choice | [1] | 4 |
| s15 size | choice | [1] | 1 |
| s16 size | choice | [1] | 1 |
| s17 size | choice | [1] | 2 |
| s18 size | choice | [1] | 4 |
| s19 size | choice | [1] | 1 |
| s20 size | choice | [1] | 1 |
| m1 size | choice | [1] | 2 |
| m2 size | choice | [1, 2] | 2 |
| m3 size | choice | [1] | 1 |
| m4 size | choice | [1, 2] | 4 |
| m5 size | choice | [1] | 1 |
| m6 size | choice | [1, 2] | 2 |
| m7 size | choice | [1] | 1 |
| m8 size | choice | [1] | 1 |
| m9 size | choice | [1] | 1 |
| m10 size | choice | [1] | 2 |
| m11 size | choice | [1] | 1 |
| m12 size | choice | [1] | 2 |
| m13 size | choice | [1] | 1 |
| m14 size | choice | [1] | 2 |
| m15 size | choice | [1] | 1 |
| m16 size | choice | [1] | 2 |
| m17 size | choice | [1] | 1 |
| m18 size | choice | [1] | 4 |
| m19 size | choice | [1] | 2 |
| m20 size | choice | [1] | 1 |
| feat dim | choice | [128, 256, 512, 1024] | 256 |

Table 11: PER CC

| Parameter | Type | Choices | Final |
|---------------|---------|--|--------------------|
| epochs | choice | [20, 30, 40] | 40 |
| lr | choice | [0.0001, 0.001, 0.01] | 0.0001 |
| momentum | uniform | [0.8, 1] | 0.8913774253680748 |
| step size | choice | [1, 2, 10] | 2 |
| gamma | uniform | [0.4, 0.8] | 0.6183381813195497 |
| act func | choice | ['ReLU', 'LeakyReLU', 'Sigmoid', 'Tanh', 'Softplus'] | ReLU |
| last act func | choice | ['Sigmoid'] | Linear |
| optimizer | choice | ['SGD', 'ADAM'] | ADAM |
| loss func | choice | ['Triplet'] | Triplet |
| branch1 | choice | [0, 1] | 0 |
| branch2 | choice | [0, 1] | 1 |
| c1 size | choice | [8, 16, 32, 64] | 16 |
| c2 size | choice | [8, 16, 32, 64] | 8 |
| c3 size | choice | [8, 16, 32, 64] | 16 |
| k1 size | choice | [1, 2, 4] | 16 |
| k2 size | choice | [1, 2, 4] | 4 |
| k3 size | choice | [1, 2, 4] | 4 |
| k4 size | choice | [1, 2, 4] | 2 |
| k5 size | choice | [1, 2, 4] | 1 |
| k6 size | choice | [1, 2, 4] | 4 |
| s1 size | choice | [1, 2] | 2 |
| s2 size | choice | [1, 2] | 2 |
| s3 size | choice | [1, 2] | 2 |
| s4 size | choice | [1, 2] | 2 |
| s5 size | choice | [1, 2] | 1 |
| s6 size | choice | [1, 2] | 2 |
| m1 size | choice | [1, 2] | 2 |
| m2 size | choice | [1, 2] | 2 |
| m3 size | choice | [1, 2] | 2 |
| m4 size | choice | [1, 2] | 1 |
| m5 size | choice | [1, 2] | 2 |
| m6 size | choice | [1, 2] | 2 |
| beta | choice | [3] | 3 |
| feat dim | choice | [128, 256, 512, 1024] | 256 |

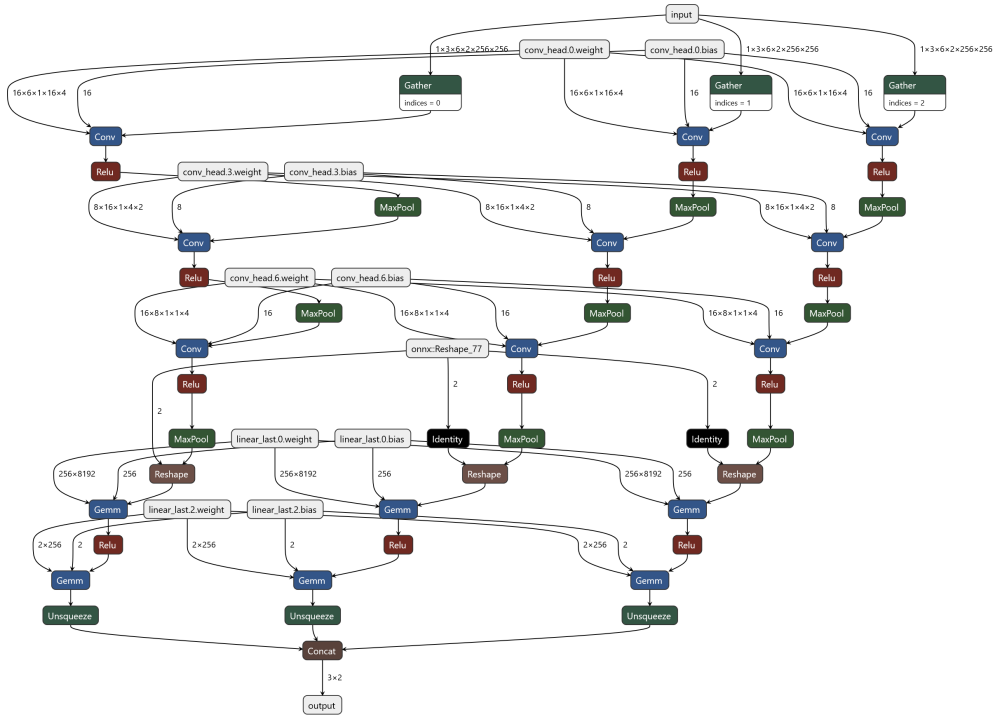


Figure 16: Structure Per CC