Institut für Informationsverarbeitung

ICML International Conference On Machine Learning

Leibniz Universität Hannover

UncertainSAM: Fast and Efficient Uncertainty Quantification of the Segment Anything Model

Timo Kaiser*

Thomas Norrenbrock

Bodo Rosenhahn

Institute for Information Processing – Leibniz University Hannover, Hannover, Germany

* Correspond to kaiser@tnt.uni-hannover.de

GitHub

Segment Anything Model (SAM)

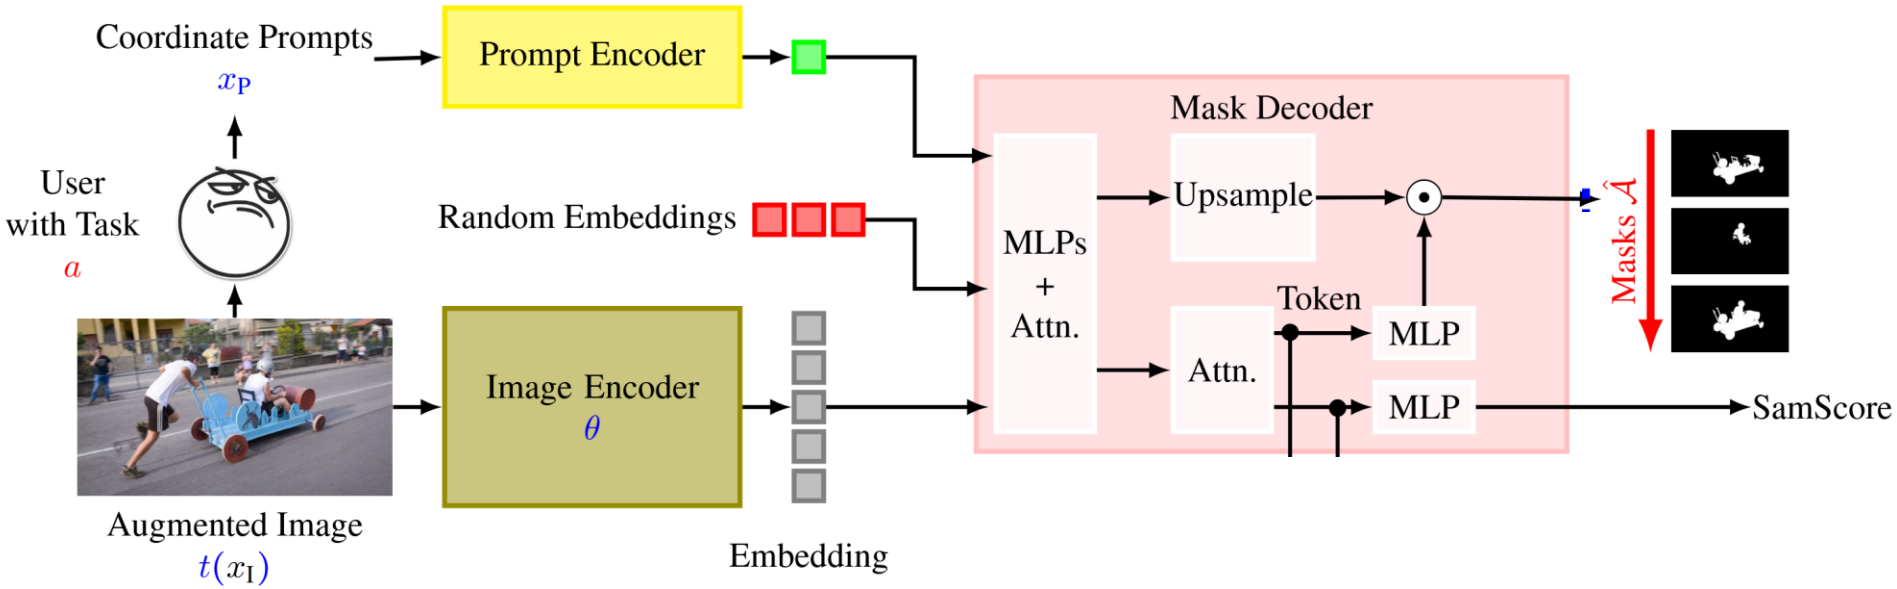

[1] Kirillov, A. et al. Segment anything.
In Proceedings of the International Conference on Computer Vision (ICCV), 2023.

Uncertainty Quantification in SAM

Uncertainty:

The likelihood, that a prediction is wrong.

Uncertainty Quantification in SAM

Uncertainty:

The likelihood, that a prediction is wrong.

What is “right” or “wrong” in Segmentation?

Intersection over Union $\frac{GT \cap Pr}{GT \cup Pr}$

Uncertainty Quantification in SAM

Uncertainty:

The likelihood, that a prediction is wrong.

What is “right” or “wrong” in Segmentation?

Intersection over Union $\frac{GT \cap Pr}{GT \cup Pr}$

Where does the Uncertainty stem from?

Model

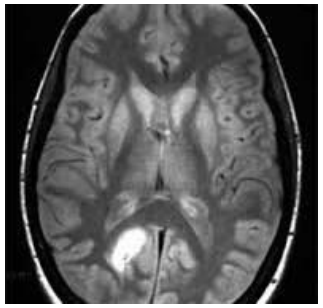

Bad Prompt

Unknown Task

UncertainSAM

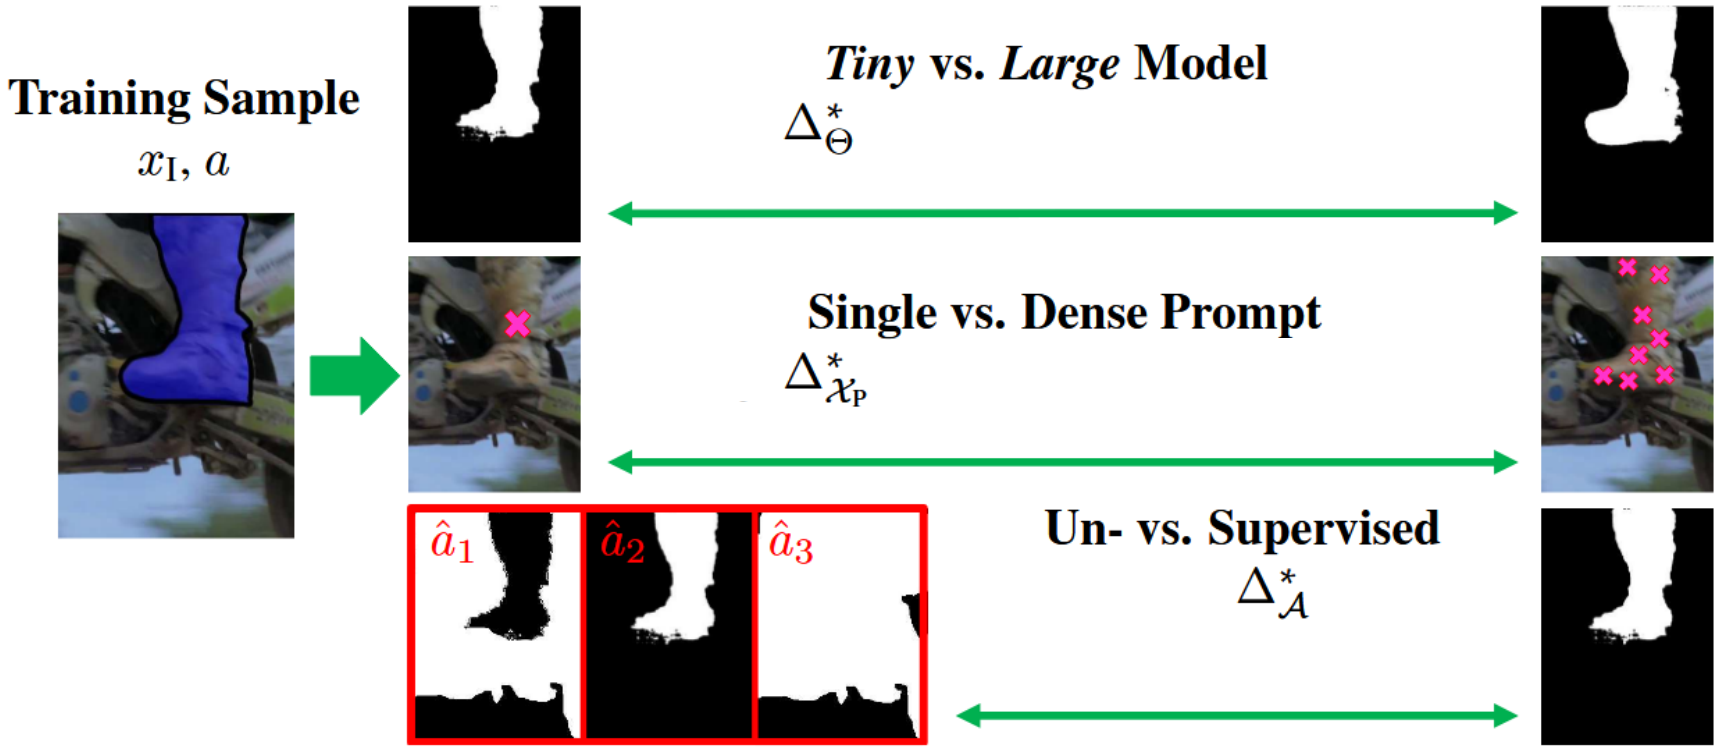

UncertainSAM

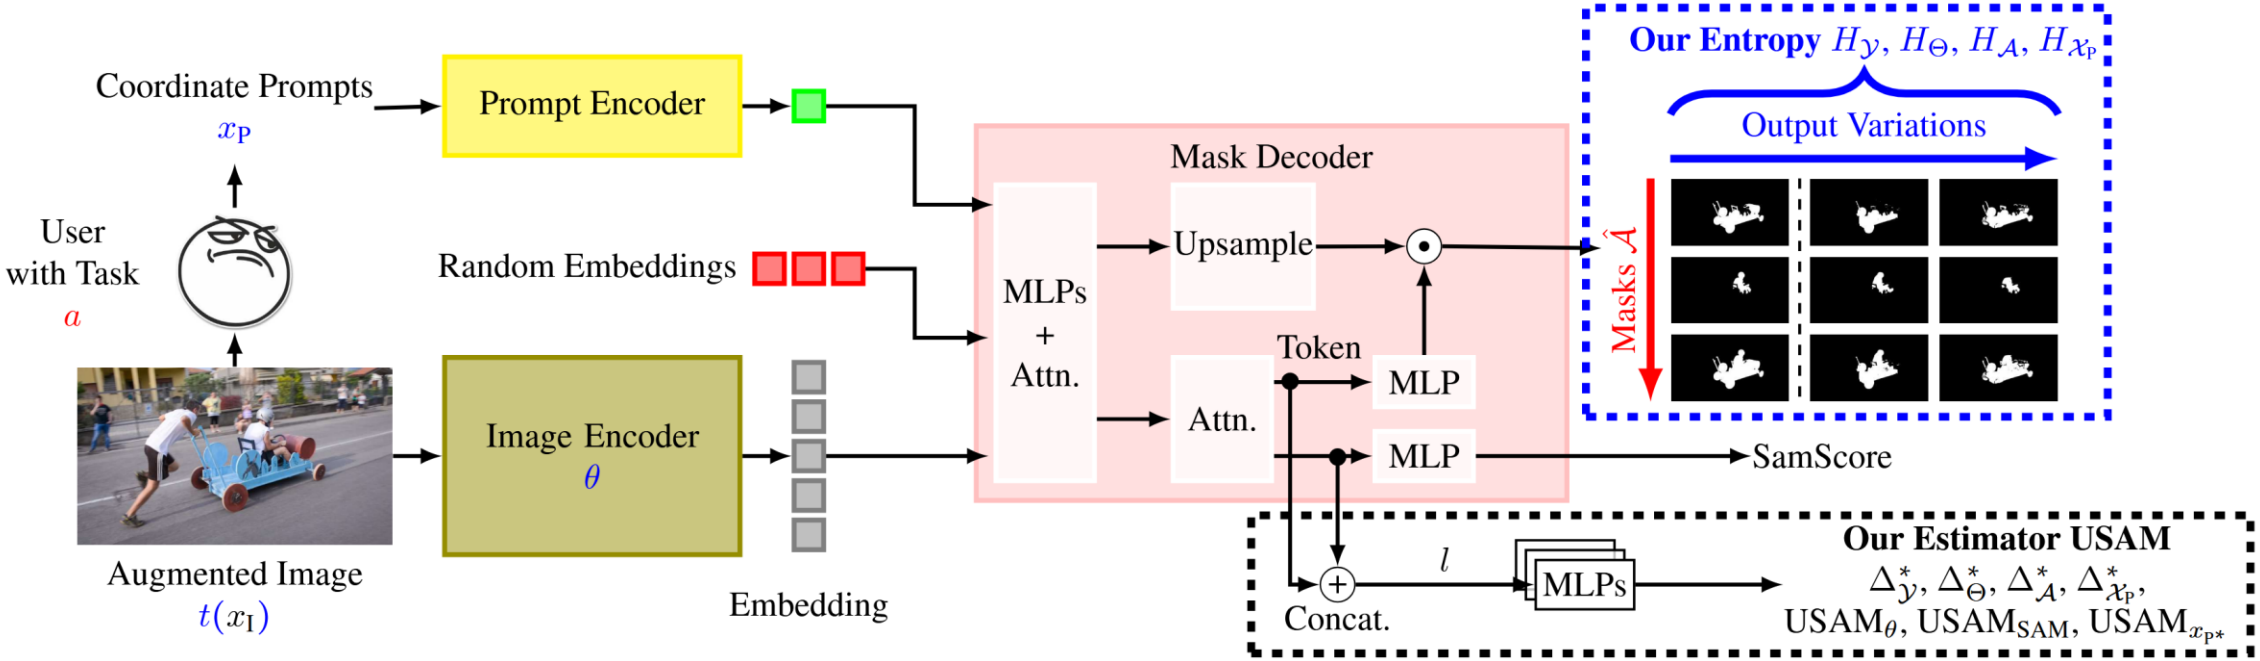

[1] Kirillov, A. et al. Segment anything.
In Proceedings of the International Conference on Computer Vision (ICCV), 2023.

Results & Insights

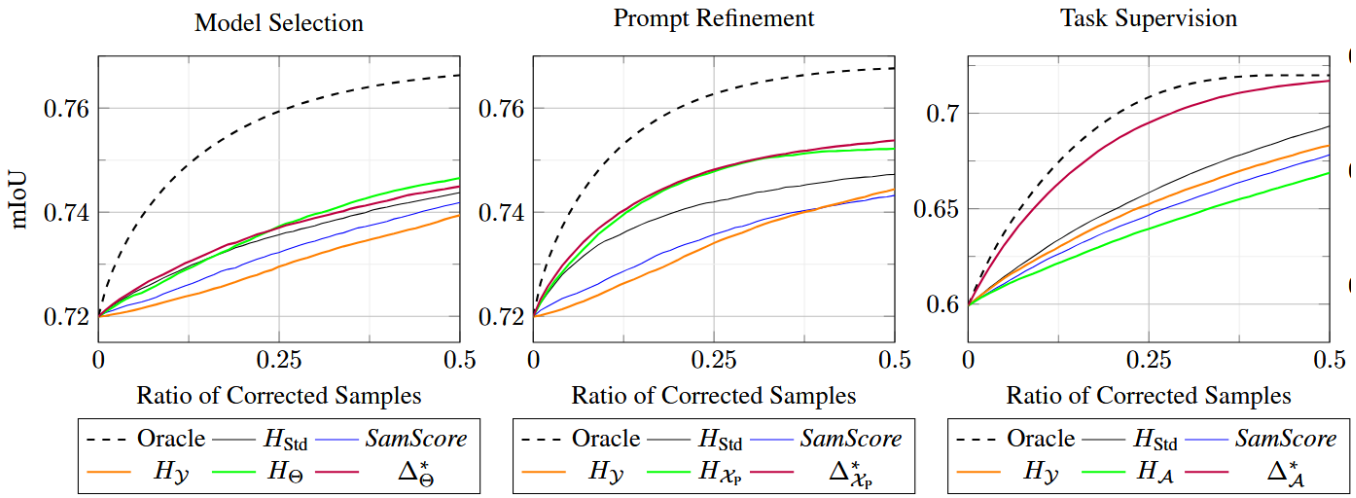

	[rel. AUC in %]↑	DAVIS	ADE20k	MOSE	COCO	SA-V
SAM	*SamScore*	64.25	52.85	68.44	57.97	58.83
	Entropy	71.78	53.48	71.56	61.61	63.49
Bayes	$H_{\mathcal{Y}}$	51.48	52.06	61.36	54.27	55.39
	H_{Θ}	**73.46**	57.65	**75.60**	64.43	**65.66**
USAM	Δ_{Θ}	66.85	60.32	71.78	63.66	61.23
	Δ^*_{Θ}	59.08	**61.55**	73.05	**66.60**	63.71

Model uncertainty quantification. Area under curve (AUC) when predicting a variable fraction of the most uncertain samples with the Large model, others with the tiny one.

	[rel. AUC in %]↑	DAVIS	ADE20k	MOSE	COCO	SA-V
SAM	*SamScore*	68.36	66.31	70.58	64.09	58.87
	Entropy	68.77	76.49	73.55	74.56	70.88
Bayes	$H_{\mathcal{Y}}$	74.88	64.53	67.32	68.12	70.50
	$H_{\mathcal{A}}$	43.86	52.79	66.04	57.55	78.13
USAM	$\Delta_{\mathcal{A}}$	94.05	**93.08**	**94.21**	94.85	94.17
	$\Delta^*_{\mathcal{A}}$	**94.31**	92.38	94.01	**94.87**	**94.61**

Task uncertainty quantification. Area under curve (AUC) when predicting a variable fraction of the most uncertain samples with the correct task, otherwise with the one selected by the SamScore.

	[rel. AUC in %] ↑	DAVIS	ADE20k	MOSE	COCO	SA-V
SAM	*SamScore*	71.82	69.12	66.85	60.55	54.13
	Entropy	77.13	70.15	69.24	68.05	63.56
Bayes	$H_{\mathcal{Y}}$	54.11	60.82	63.74	61.25	55.78
	$H_{\mathcal{X}_P}$	**80.75**	79.41	74.33	74.20	67.69
USAM	$\Delta_{\mathcal{X}_P}$	75.04	83.23	74.21	78.17	70.15
	$\Delta^*_{\mathcal{X}_P}$	75.53	**83.41**	**74.50**	**78.89**	**71.27**

Prompt uncertainty quantification. Area under curve (AUC) when predicting a variable fraction of the most uncertain samples with a refined prompt containing multiple point coordinates, others with a single-coordinate prompt.

Results & Insights

$\left[\frac{\text{Seconds}}{\text{Iteration}}\right] \downarrow$	SAM	+Entropy	$+\lvert\mathcal{T}\rvert = 5$	$+\lvert\mathcal{X}_\text{P}\rvert = 8$	+USAM
Large	0.437	0.452	2.187	0.500	**0.441**
Base+	0.205	0.233	1.028	0.289	**0.210**
Small	0.134	0.157	0.688	0.232	**0.142**
Tiny	0.122	0.149	0.584	0.198	**0.139**

Runtime of SAM with and without UQ methods on a regular image performed on a NVIDIA RTX3050 Ti.

Mask Token	IoU Token	Model Uncertainty	Prompt Uncertainty	Task Uncertainty
✗	✓	61.19%	72.08%	86.89%
✓	✗	62.63%	76.42%	91.96%
✓	✓	**63.66%**	**78.30%**	**94.82%**

Token ablation. The UQ performance of USAM when removing mask or IoU tokens from the MLP input on the COCO dataset, measured in relative AUC as in the main experiments.

Samples

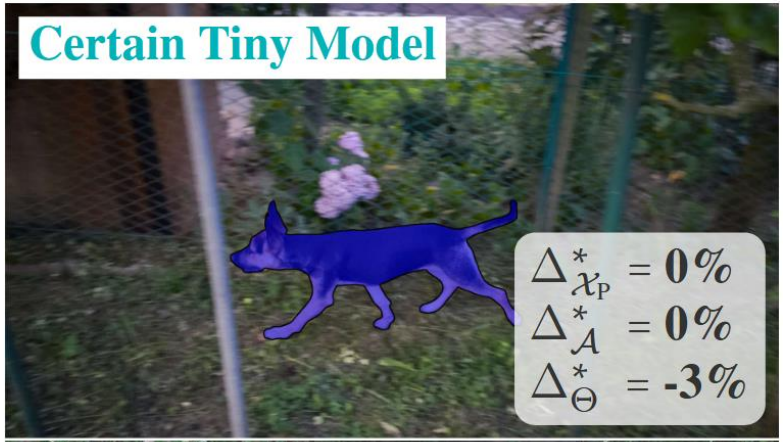

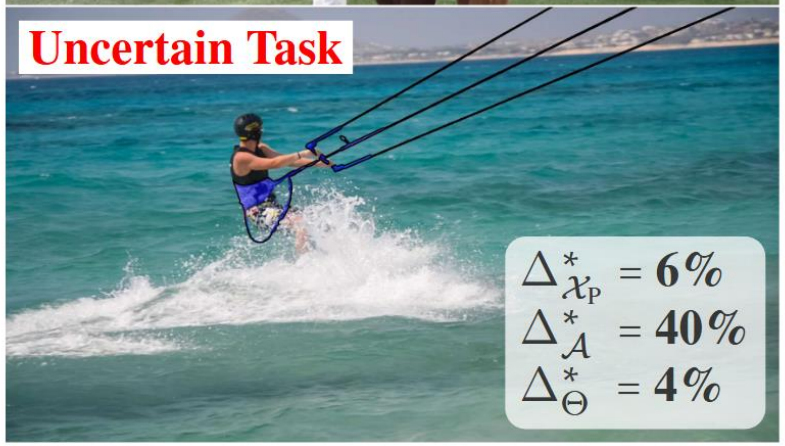

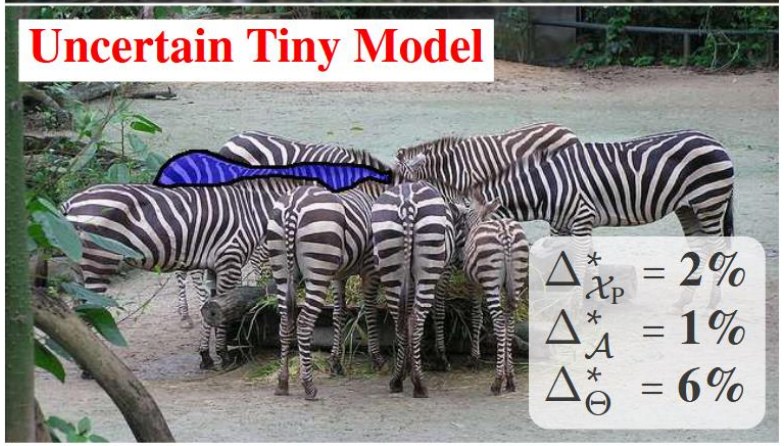

$\Delta^*_{X_P}$: Estimated Prompt Gap Δ^*_{A}: Estimated Task Gap Δ^*_{Θ}: Estimated Model Gap

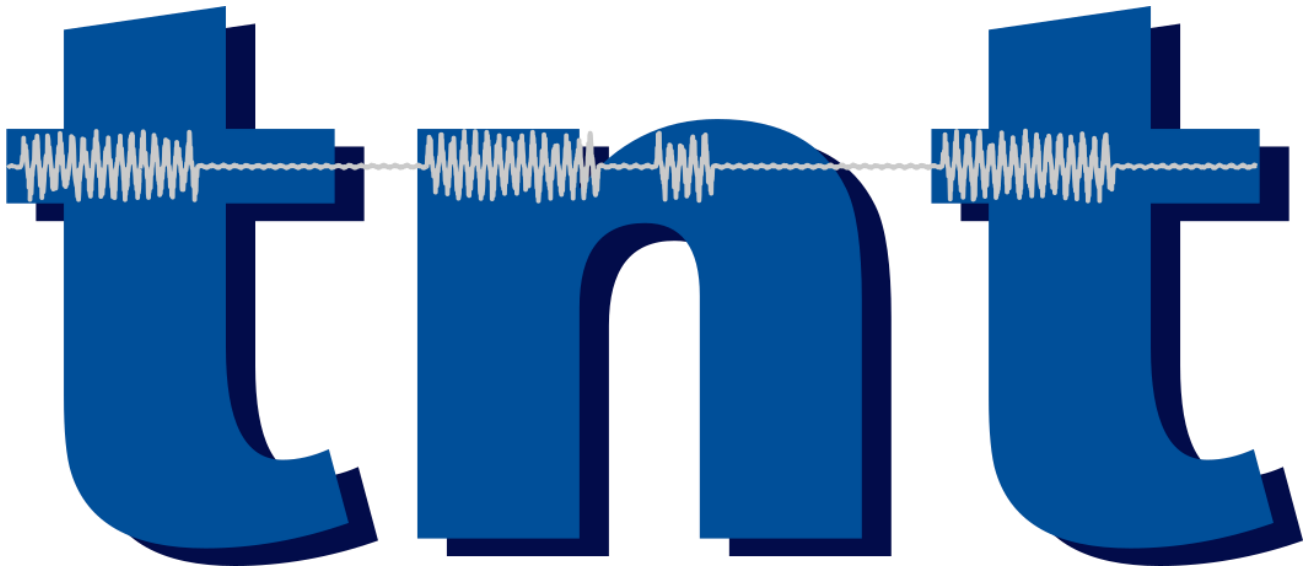

Questions and discussion are welcome! During the Poster Session or at

kaiser@tnt.uni-hannover.de or +49 511-762-19504

GitHub