DAGM German Conference on Pattern Recognition, Freiburg, September 23 - September 26, 2025

The German Conference on Pattern Recognition (GCPR) is the annual symposium of the German Association for Pattern Recognition (DAGM). It is the national venue for recent advances in image processing, pattern recognition, and computer vision, and it follows the long tradition of the DAGM conference series, which has been renamed to GCPR in 2013 to reflect its increasing internationalization. In 2025, the 47th conference will be hosted by the University of Freiburg.

Conference Venue

The venue for DAGM-GCPR 2025 is inside the prestigious university building KG1, Platz der Universität 3 in Freiburg. The talks will be given in room 1115 (Aula), posters are presented in 1114 (Vorraum Aula) and in the entrance hall (Prometheus Halle).

Program at a glance

Date Time	September 23	September 24	September 25	September 26
9:00		Welcome & Awards session (1115)	Poster Session 2 (1114 and PH)	**Keynote**: **Efstratios Gavves** (1115)
9:30				
10:00			Coffee	Oral Session 5 (1115)
10:30		Coffee	**Keynote**: **Dima Damen** Opportunities in Egocentric Vision (1115)	Coffee
11:00		**Keynote**: **Venkatesh Babu Radhakrishnan** Towards Fair and Controllable Diffusion Models (1115)		
11:30			Oral Session 4 (1115)	Poster Session 4 (1114 and PH)
12:00		Oral Session 1 (1115)		
12:30		Lunch	Lunch	Goodbye
13:00	Tutorial: AI Agents by Peter Gehler and Matthias Kümmerer (1115)			
13:30		Oral Session 2 (1115)	Poster Session 3 (1114 and PH)	
14:00				
14:30	TK Meeting	**Keynote**: **Alex Kolesnikov** A Journey Toward Unified Vision Models (1115)		
15:00			Coffee	
15:30		Coffee	**Keynote**: **Stefanie Jegelka** (1115)	
16:00		Poster Session 1 (1114 and PH)		
16:30			Oral session	
17:00		Oral session 3 (1115)	Get to the City tour City Tour	
17:30				
18:00	Industry Meetup & Welcome Reception (1115, 1114 and PH)	DAGM Assembly (1115)		
18:30			Dinner	

Tuesday 23.9.

13:00-17:30 Tutorial: AI Agents

Matthias Kümmerer, Linus Schneider, Jaisidh Singh, Robin Ruff, Mamen Chembakasseril, Peter Gehler

Always wanted to learn what AI Agents are about, but never got the time? Join our half-day tutorial on LangGraph, LangChain, and the Model Context Protocol (MCP)—three puzzle pieces that snap together to turn large-language-model theory into fully fledged, experiment-ready agents. We open with explanations of all software packages and their concepts, and have some examples. Then we dive straight into hands-on hacking, where you will write your own AI agent or MCP host.

Bring a laptop, your favorite IDE, and an LLM key (we can provide an LLM key). If you have an idea about what you would like to code, bring that as well; otherwise, we have some examples to work on together. This tutorial is both instructional and hands-on, your chance to get started with agentic frameworks. Let's learn together

13:00-15:00 Presentation

15:00-15:30 Coffee Break

15:30-17:30 Hands-on

.

17:30 Welcome Reception and Industry Meet-up

The Welcome Reception is sponsored by Zeiss, and the Industry Meet-up is organized and sponsored by the Tübingen AI Center.

17:30 Entry

18:00 Impulse presentation by SPRIN-D

18:20 Company pitches

19:00 Welcome Reception with Industry posters

Participation from Carl Zeiss, MVTec, Zalando, neura-robotics, Bosch, scholar-inbox, KI macht Schule, KI Allianz BW, Deep Scenario, Viscoda, Zebracat AI

Along with the welcome reception of the conference we host an industry meet-up, where local companies, sponsors, and academia can meet and talk. Many German vision professors, PhD and Master students, and researchers will be present. There will be plenty of time to talk with a glass of wine, a bottle of beer, or a bottle of lemonade. We expect 150-200 people.

If you are a local company and you are interested to join, please contact the industry chair.

17:30 Entry

18:00 Impulse presentation by SPRIN-D

18:20 Company pitches

Wednesday 24.9.

9:00 Opening and Awards

9:00 Opening
9:15 Award talk
10:00 Award talk

10:30 Coffee break

11:00 Keynote: Venkatesh Babu Radhakrishnan - Towards Fair and Controllable Diffusion Models

Abstract: Diffusion models have transformed text-to-image generation, but challenges remain in fairness, representativeness, and user control. In this talk, we present some of our efforts that address these critical gaps. We begin by examining the demographic and geographic biases in popular generative models, showing over-representation of certain regions and attributes. To mitigate the biases in generative models, we propose distribution-guided debiasing methods that align outputs with desired attribute distributions without retraining, enabling fairer and more inclusive generations. Beyond fairness, we introduce fine-grained control mechanisms, enabling precise attribute editing and identity preservation, bridging realism with user-driven customization. We extend controllability to spatial reasoning with affordance-aware text-guided human placement, ensuring semantically plausible compositions, while the proposed zero-shot, depth-aware editing enables realistic scene modifications without additional supervision. We hope these contributions help in making the generative models that are equitable, transparent, and highly controllable for real-world applications.

12:00 Oral Session 1

- VisualChef: Generating Visual Aids in Cooking via Mask Inpainting
  Kuzyk, Oleh; Li, Zuoyue; Pollefeys, Marc; Wang, Xi
- Investigating Structural Pruning and Recovery Techniques for Compressing Multimodal Large Language Models: An Empirical Study
  Huang, Yiran; Thede, Lukas; Mancini, Massimiliano; Xu, Wenjia; Akata, Zeynep

12:30 Lunch break

University Cafeteria Rempartstraße 18 (opposite to the conference site). Bring the voucher from your registration package!

13:30 Oral Session 2

- Assessing Foundation Models for Mold Colony Detection with Limited Training Data
  Pichler, Henrik; Keuper, Janis; Copping, Matthew
- Out-of-Distribution Detection in LiDAR Semantic Segmentation Using Epistemic Uncertainty from Hierarchical GMMs
  Shojaei Miandashti, Hanieh; Brenner, Claus
- LADB: Latent Aligned Diffusion Bridges for Semi-Supervised Domain Translation
  Wang, Xuqin; Wu, Tao; Zhang, Yanfeng; Liu, Lu; Wang, Dong; Sun, Mingwei; Wang, Yongliang; Zeller, Niclas; Cremers, Daniel
- Activation Subspaces for Out-of-Distribution Detection
  Zöngür, Barış; Hesse, Robin; Roth, Stefan

14:30 Keynote: Alex Kolesnikov - A Journey Toward Unified Vision Models

Abstract: Text-only transformer models thrive on a single, simple interface: next-token prediction. They then let scale and data do the heavy lifting. By contrast, models in the vision domain remain fragmented and are hindered by non-trivial components such as box proposals, non-maximum suppression, and matching losses. In this talk, I'll share insights from my research journey toward simpler, more unified vision models.

I'll trace a path through three projects. First, UViM shows how to express structured outputs, such as panoptic segmentation masks and depth maps, as discrete codes that a vanilla autoregressive transformer can generate. Next, I'll dive into policy-gradient RL fine-tuning, addressing fundamental limitations of pure log-likelihood training by optimizing the metrics we actually care about. Finally, I'll introduce JetFormer, a decoder-only autoregressive transformer capable of full end-to-end modeling of high-resolution images.

15:30 Poster Session 1 and Coffee break (Coffee until 16:00)

P1	VisualChef: Generating Visual Aids in Cooking via Mask Inpainting	Kuzyk, Oleh; Li, Zuoyue; Pollefeys, Marc; Wang, Xi
P2	Investigating Structural Pruning and Recovery Techniques for Compressing Multimodal Large Language Models: An Empirical Study	Huang, Yiran; Thede, Lukas; Mancini, Massimiliano; Xu, Wenjia; Akata, Zeynep
P3	EVCS: A Benchmark for Fine-Grained Electric Vehicle Charging Station Detection	Chen, Lin; Südbeck, Sönke; Riggers, Christoph; Geib, Tobias; Cordes, Kai; Broszio, Hellward
P4	NaT-ReX: Naturalness Assessment with Transformer-Based Reliable Explainability	Emam, Ahmed; Farag, Mohamed; Russwurm, Marc; Roscher, Ribana
P5	Assessing Foundation Models for Mold Colony Detection with Limited Training Data	Pichler, Henrik; Keuper, Janis; Copping, Matthew
P6	Out-of-Distribution Detection in LiDAR Semantic Segmentation Using Epistemic Uncertainty from Hierarchical GMMs	Shojaei Miandashti, Hanieh; Brenner, Claus

P7	subCellSAM: Zero-Shot (Sub-)Cellular Segmentation for Hit Validation in Drug Discovery	Hanimann, Jacob; Siegismund, Daniel; Wieser, Mario; Steigele, Stephan
P8	Efficient Masked Attention Transformer for Few-Shot Classification and Segmentation	Carrion, Dustin; Roth, Stefan; Schaub-Meyer, Simone
P9	Object Risk Estimation for Autonomous Driving Safety	Khan, Abdul Hannan; Shafiq, Syed; van Elst, Ludger; Dengel, Andreas
P10	LADB: Latent Aligned Diffusion Bridges for Semi-Supervised Domain Translation	Wang, Xuqin; Wu, Tao; Zhang, Yanfeng; Liu, Lu; Wang, Dong; Sun, Mingwei; Wang, Yongliang; Zeller, Niclas; Cremers, Daniel
P11	Common Data Properties Limit Object-Attribute Binding in CLIP	Gurung, Bijay; Hoffmann, David; Brox, Thomas
P12	Activation Subspaces for Out-of-Distribution Detection	Zöngür, Barış; Hesse, Robin; Roth, Stefan
P13	SegSLR: Promptable Video Segmentation for Isolated Sign Language Recognition	Schreiber, Sven; Sarhan, Noha; Frintrop, Simone; Wilms, Christian
P14	MT-Occ: Single-View 3D Occupancy Prediction via Multi-Task Distillation	Li, Zhi; Aljundi, Rahaf; Reino, Daniel; Schiele, Bernt
P16 (PH)	FedPCE: Federated Personalized Client Embeddings for Post-training Knowledge Distillation	Hansel, Soma; Kobler, Erich; Effland, Alexande
P18 (PH)	Rethinking Semi-supervised Segmentation Beyond Accuracy: Robustness and Reliability	Landgraf, Steven; Hillemann, Markus; Ulrich, Markus
P20 (PH)	A Cascaded Dilated Convolution Approach for Mpox Lesion Classification	Deshmukh, Ayush
P22 (PH)	Detection of Synthetic Face Images: Accuracy, Robustness, Generalization	Petrželková, Nela; Čech, Jan
PH = Prometheus Halle		

17:00 Oral Session 3

- Video Object Segmentation-aware Audio Generation
  Viertola, Ilpo; Iashin, Vladimir; Rahtu, Esa
- Combining Absolute and Semi-Generalized Relative Poses for Visual Localization
  Panek, Vojtech; Sattler, Torsten; Kukelova, Zuzana
- sshELF: Single-Shot Hierarchical Extrapolation of Latent Features for 3D Reconstruction from Sparse-Views
  Najafli, Eyvaz; Kästingschäfer, Marius; Bernhard, Sebastian; Brox, Thomas; Geiger, Andreas
- VGGSounder: Audio-Visual Evaluations for Foundation Models
  Zverev, Daniil; Wiedemer, Thaddäus; Prabhu, Ameya; Bethge, Matthias; Brendel, Wieland; Koepke, A. Sophia

18:00 DAGM Assembly Only for DAGM members. Become a member.

Thursday 25.9.

9:00 Poster Session 2 (Coffee from 10:00)

P1	Video Object Segmentation-aware Audio Generation	Viertola, Ilpo; Iashin, Vladimir; Rahtu, Esa
P2	Combining Absolute and Semi-Generalized Relative Poses for Visual Localization	Panek, Vojtech; Sattler, Torsten; Kukelova, Zuzana
P3	synth-dacl: Does Synthetic Defect Data Enhance Segmentation Accuracy and Robustness for Real-World Bridge Inspections?	Flotzinger, Johannes*; Deuser, Fabian; Jaziri, Achref ; Neumann, Heiko; Oswald, Norbert; Ramesh, Visvanathan; Braml, Thomas
P4	Hierarchical Insights: Exploiting Structural Similarities for Reliable 3D Semantic Segmentation	Dreissig, Mariella; Ruehle, Simon; Piewak, Florian; Boedecker, Joschka
P5	Semantic Segmentation of Structural Damage: A Comparative Study of YOLO11 and Encoder-Decoder Networks	Krefft, Lorenz; Hoegner, Ludwig
P6	HistDiST: Histopathological Diffusion-based Stain Transfer	Grosskopf, Erik; Bundele, Valay; Hosseinzadeh, Mehran; Lensch, Hendrik
P7	Structured Universal Adversarial Attacks on Object Detection for Video Sequences	Jacob, Sven; Shao, Weijia; Kasneci, Gjergji
P8	Unlocking In-Context Learning for Natural Datasets Beyond Language Modelling	Bratulić, Jelena; Mittal, Sudhanshu; Hoffmann, David; Böhm, Samuel; Schirrmeister, Robin; Ball, Tonio; Rupprecht, Christian; Brox, Thomas
P9	Graph Roof Reconstruction with Synthetic Data from Misaligned Labels	Amrullah, Chaikal; Bittner, Ksenia
P10	sshELF: Single-Shot Hierarchical Extrapolation of Latent Features for 3D Reconstruction from Sparse-Views	Najafli, Eyvaz; Kästingschäfer, Marius; Bernhard, Sebastian; Brox, Thomas; Geiger, Andreas
P11	VGGSounder: Audio-Visual Evaluations for Foundation Models	Zverev, Daniil; Wiedemer, Thaddäus; Prabhu, Ameya; Bethge, Matthias; Brendel, Wieland; Koepke, A. Sophia
P12	CausalRivers - Scaling up benchmarking of causal discovery for real-world time-series	Stein, Gideon; Shadaydeh, Maha; Blunk, Jan; Penzel, Niklas; Denzler, Joachim
P13	Can Multitask Learning Enhance Model Explainability?	Najjar, Hiba; Alshbib, Bushra; Dengel, Andreas
P14	DASH: Detection and Assessment of Systematic Hallucinations of VLMs	Augustin, Maximilian; Neuhaus, Yannic; Hein, Matthias
P16 (PH)	DCBM: Data-Efficient Visual Concept Bottleneck Models	Prasse, Katharina; Knab, Patrick; Marton, Sascha; Bartelt, Christian;

		Keuper, Margret
P18 (PH)	Higher-Order Ratio Cycles for Fast and Globally Optimal Shape Matching	Roetzer, Paul; Ehm, Viktoria; Cremers, Daniel; Lähner, Zorah; Bernard, Florian
P20 (PH)	PhysicsGen: Can Generative Models Learn from Images to Predict Complex Physical Relations?	Spitznagel, Martin; Vaillant, Jan; Keuper, Janis
P22 (PH)	High-Resolution 3D Shape Matching with Global Optimality and Geometric Consistency	El Amrani, Nafie; Rötzer, Paul; Bernard, Florian
P23 (PH)	Faster Inference of Flow-Based Generative Models via Improved Data-Noise Coupling	Davtyan, Aram; Dadi, Leello Tadesse; Cevher, Volkan; Favaro, Paolo

10:30 Keynote: Dima Damen - Opportunities in Egocentric Vision

Abstract: Forecasting the rise of wearable devices equipped with audio-visual feeds, this talk will present opportunities for research in egocentric video understanding. The talk argues for new ways to foresee egocentric videos as partial observations of a dynamic 3D world, where objects are out of sight but not out of mind. I'll review new data collection and annotation HD-EPIC (https://hd-epic.github.io/) that merges video understanding with 3D modelling, showcasing current failures of VLMs in understanding the perspective outside the camera's field of view — a task trivial for humans.

11:30 Oral Session 4

- MCUCoder: Adaptive Bitrate Learned Video Compression for IoT Devices
  Hojjat, Ali; Haberer, Janek; Landsiedel, Olaf
- CoProU-VO: Combining Projected Uncertainty for End-to-End Unsupervised Monocular Visual Odometry
  Xie, Jingchao; Dhaouadi, Oussema; Chen, Weirong; Meier, Johannes; Kaiser, Jacques; Cremers, Daniel
- Deep Learning-Assisted Dynamic Mode Decomposition for Non-resonant Background Removal in CARS Spectroscopy
  Chalain Valapil, Adithya Ashok; Messerschmidt, Carl; Shadaydeh, Maha; Schmitt, Michael; Popp, Jürgen; Denzler, Joachim
- Spatial Reasoning with Denoising Models
  Wewer, Christopher; Pogodzinski, Bartlomiej; Schiele, Bernt; Lenssen, Jan

12:30 Lunch break

University Cafeteria Rempartstraße 18 (opposite to the conference site). Bring the voucher from your registration package!

13:30 **Poster Session 3** (Coffee from 15:00)

P01	MCUCoder: Adaptive Bitrate Learned Video Compression for IoT Devices	Hojjat, Ali; Haberer, Janek; Landsiedel, Olaf
P02	CoProU-VO: Combining Projected Uncertainty for End-to-End Unsupervised Monocular Visual Odometry	Xie, Jingchao; Dhaouadi, Oussema; Chen, Weirong; Meier, Johannes; Kaiser, Jacques; Cremers, Daniel
P03	Banded Square Root Matrix Factorization for Differentially Private Model Training	Kalinin, Nikita; Lampert, Christoph
P04	Do It Yourself: Learning Semantic Correspondence from Pseudo-Labels	Dünkel, Olaf; Wimmer, Thomas; Theobalt, Christian; Rupprecht, Christian; Kortylewski, Adam
P05	γ-Quant: Towards Learnable Quantization for Low-bit Pattern Recognition	Fatima, Mishal; Agnihotri, Shashank; Bock, Marius; Gandikota, Kanchana; Van Laerhoven , Kristof; Moeller, Michael; Keuper, Margret
P06	Deep Learning-Assisted Dynamic Mode Decomposition for Non-resonant Background Removal in CARS Spectroscopy	Chalain Valapil, Adithya Ashok; Messerschmidt, Carl; Shadaydeh, Maha; Schmitt, Michael; Popp, Jürgen; Denzler, Joachim
P07	Spatial Reasoning with Denoising Models	Wewer, Christopher; Pogodzinski, Bartlomiej; Schiele, Bernt; Lenssen, Jan
P08	The GOOSE Dataset for Perception in Unstructured Environments	Mortimer, Peter; Hagmanns, Raphael; Granero, Miguel; Petereit, Janko; Luettel, Thorsten
P09	Can Multitask Learning Enhance Model Explainability?	Najjar, Hiba; Alshbib, Bushra; Dengel, Andreas
P10	RadarSeq: A Temporal Vision Framework for User Churn Prediction via Radar Sequence Chart	Najafi, Sina; Sepanj, M.Hadi; Jafari, Fahimeh
P11	FlowBench: Benchmarking Optical Flow Estimation Methods for Reliability and Generalization	Agnihotri, Shashank; Caspary, Julian; Schwarz, Luca; Gao, Xinyan; Schmalfuss, Jenny; Bruhn, Andres; Keuper, Margret
P12	VITAL: More Understandable Feature Visualization through Distribution Alignment and Relevant Information Flow	Ada Görgün; Bernt Schiele; Jonas Fischer
P13	B-cosification: Transforming Deep Neural Networks to be Inherently Interpretable	Arya, Shreyash; Rao, Sukrut; Böhle, Moritz; Schiele, Bernt
P14	CutS3D: Cutting Semantics in 3D for 2D Unsupervised Instance Segmentation	Sick, Leon; Engel, Dominik; Hartwig, Sebastian; Hermosilla, Pedro; Ropinski, Timo
P15 (PH)	Scribbles for All: Benchmarking Scribble Supervised Segmentation Across Datasets	Boettcher, Wolfgang; Hoyer, Lukas; Uenal, Ozan; Lenssen, Jan Eric; Schiele, Bernt
P16 (PH)	TikZero: Zero-Shot Text-Guided Graphics Program Synthesis	Belouadi, Jonas; Ilg, Eddy; Keuper, Margret; Tanaka, Hideki; Utiyama,

		Masao; Dabre, Raj; Eger, Steffen Eger; Ponzetto, Simone Paolo
P17 (PH)	LeGrad: An Explainability Method for Vision Transformers via Feature Formation Sensitivity	Bousselham, Walid; Boggust, Angie; Chaybouti, Sofian; Strobelt, Hendrik; Kühne, Hilde
P18 (PH)	Using Shapley interactions to understand how models use structure	Divyansh Singhvi, Diganta Misra, Andrej Erkelens, Raghav Jain, Isabel Papadimitriou, Naomi Saphra
P19 (PH)	Scene-Centric Unsupervised Panoptic Segmentation	Hahn, Oliver; Reich, Christoph; Araslanov, Nikita; Cremers, Daniel; Rupprecht, Christian; Roth, Stefan
P20 (PH)	MANTA: Diffusion Mamba for Efficient and Effective Stochastic Long-Term Dense Action Anticipation	Zatsarynna, Olga; Bahrami, Emad; Abu Farha, Yazan; Francesca, Gianpiero; Gall, Jürgen
P21 (PH)	HydraViT: Stacking Heads for a Scalable ViT	Haberer, Janek; Hojjat, Ali; Landsiedel, Olaf
P22 (PH)	CAGE: Unsupervised Visual Composition and Animation for Controllable Video Generation	Davtyan, Aram; Sameni, Sepehr; Ommer, Björn; Favaro, Paolo
P23 (PH)	Can LLMs Separate Instructions From Data? And What Do We Even Mean By That?	Sahar Abdelnabi; Soroush Tabesh; Mario Fritz; Christoph Lampert

15:30 Keynote: Stefanie Jegelka - Title: Does computational structure tell us about deep learning? Some thoughts and examples

Abstract: Understanding and steering deep learning training and inference is a nontrivial endeavor. In this talk, I will look at training, learning, and inference from the perspective of computational structure, via a few diverse examples.

First, computational structure may help understand expressiveness and biases in deep learning models. For instance, it can connect graph neural networks to SDPs, indicating their capability of learning optimal approximation algorithms. It can also help explain position biases in LLMs. Second, computational structure exists not only in the architecture but also in inference procedures such as the chain-of-thought. Finally, if time permits, we will connect the architectural structure via neural parameter symmetries to the training and loss landscape of deep models and explore the effect of removing symmetries.

16:30 Oral Session

- Unlocking In-Context Learning for Natural Datasets Beyond Language Modelling, Bratulić, Jelena; Mittal, Sudhanshu; Hoffmann, David; Böhm, Samuel; Schirrmeister, Robin; Ball, Tonio; Rupprecht, Christian; Brox, Thomas

17:00 Guided City Tour. The meeting point will be at "Platz der Alten Synagoge", just next to the conference venue.

18:30 Conference Dinner

The conference dinner is sponsored by **Black Forest Labs**.

It is at Dattler Schlossbergrestaurant on Schlossberg with a nice view over the city of Freiburg. You can either walk up the hill to the restaurant or you can take the cog railway from Stadtgarten. See the venue for the map. You must bring your badge.

Friday 26.9.

9:00 Keynote: Efstratios Gavves - Title: TBD

10:00 Oral Session 5

- Using Knowledge Graphs to harvest datasets for efficient CLIP model training
  Ging, Simon; Walter, Sebastian; Bratulić, Jelena; Dienert, Johannes; Bast, Hannah; Brox, Thomas
- Feed-Forward SceneDINO for Unsupervised Semantic Scene Completion
  Jevtić, Aleksandar; Reich, Christoph; Wimbauer, Felix; Hahn, Oliver; Rupprecht, Christian; Roth, Stefan; Cremers, Daniel

10:30 Poster Session 4 and Coffee break (Coffee until 11:00)

P01	Using Knowledge Graphs to harvest datasets for efficient CLIP model training	Ging, Simon; Walter, Sebastian; Bratulić, Jelena; Dienert, Johannes; Bast, Hannah; Brox, Thomas
P02	Box it and Track it: A Weakly Supervised Framework for Cell Tracking	Khalid, Nabeel; Koochali, Mohammadmahdi; Naseem, Khola; Lovell, Gillian; Migliori, Bianca; Porto, Daniel; Trygg, Johan; Dengel, Andreas; Ahmed, Sheraz
P03	AIM: Amending Inherent Interpretability via Self-Supervised Masking	Alshami, Eyad; Agnihotri, Shashank; Schiele, Bernt; Keuper, Margret
P04	Towards Optimizing Large-Scale Multi-Graph Matching in Bioimaging	Kahl, Max; Stricker, Sebastian; Hutschenreither, Lisa; Bernard, Florian; Rother, Carsten; Savchynskyy, Bogdan
P05	StorySync: Training-Free Subject Consistency via Region Harmonization	Gaur, Gopalji; Zolfaghari, Mohammadreza; Brox, Thomas
P06	Don’t Miss Out on Novelty: Importance of Novel Features for Deep Anomaly Detection	Sivaprasad, Sarath; Fritz, Mario
P07	When and How Does CLIP Enable Domain and Compositional Generalization?	Kempf, Elias; Schrodi, Simon; Argus, Max; Brox, Thomas
P08	Two Effects, One Trigger: On the Modality Gap, Object Bias, and Information Imbalance in Contrastive Vision-Language Models	Schrodi, Simon; Hoffmann, David; Argus, Max; Fischer, Volker; Brox, Thomas
P09	On the Dangers of Bootstrapping Generation for Continual Learning and Beyond	Zverev, Daniil; Koepke, Almut Sophia; Henriques, Joao
P10	Road Obstacle Video Segmentation	Rai, Shyam Nandan; Karthik, Shyamgopal; Georgescu, Iuliana; Caputo, Barbara; Masone, Carlo ; Akata, Zeynep

P11	SOS: Segment Object System for Open-World Instance Segmentation With Object Priors	Wilms, Christian; Rolff, Tim; Hillemann, Maris; Johanson, Robert; Frintrop, Simone
P12	Feed-Forward SceneDINO for Unsupervised Semantic Scene Completion	Jevtić, Aleksandar; Reich, Christoph; Wimbauer, Felix; Hahn, Oliver; Rupprecht, Christian; Roth, Stefan; Cremers, Daniel
P13	Can LLMs Separate Instructions From Data? And What Do We Even Mean By That?	Zverev, Egor; Abdelnabi, Sahar; Tabesh, Soroush; Fritz, Mario ; Lampert, Christoph
P14	Removing Cost Volumes from Optical Flow Estimators	Kiefhaber, Simon; Roth, Stefan; Schaub-Meyer, Simone
P15 (PH)	Combined Image Data Augmentations diminish the benefits of Adaptive Label Smoothing	Siedel, Georg; Gupta, Ekagra; Shao, Weijia; Vock, Silvia; Morozov, Andrey
P16 (PH)	Can We Talk Models Into Seeing the World Differently?	Gavrikov, Paul; Keuper, Janis
P17 (PH)	FAIR-TAT: Improving Model Fairness Using Targeted Adversarial Training	Medi, Tejaswini
P18 (PH)	FastCAV: Efficient Computation of Concept Activation Vectors for Explaining Deep Neural Networks	Schmalwasser, Laines; Penzel, Niklas; Denzler, Joachim; Niebling, Julia
P19 (PH)	Probabilistic Embeddings for Frozen Vision-Language Models: Uncertainty Quantification with Gaussian Process Latent Variable Models	Venkataramanan, Aishwarya; Bodesheim, Paul; Denzler, Joachim
P20 (PH)	SEED4D: A Synthetic Ego–Exo Dynamic 4D Data Generator, Driving Dataset and Benchmark	Kästingschäfer, Marius; Gieruc, Theo; Bernhard, Sebastian; Campbell, Dylan; Insafutdinov, Eldar; Najafli, Eyvaz; Brox, Thomas
P21 (PH)	Electromyography-Informed Facial Expression Reconstruction for Physiological-Based Synthesis and Analysis	Büchner, Tim; Anders, Christoph; Guntinas-Lichius, Orlando; Denzler, Joachim
P22 (PH)	Implicit Language Models are RNNs: Balancing Parallelization and Expressivity	Schoene, Mark; Rahmani, Babak; Kremer, Heiner; Falck, Fabian; Ballani, Hitesh; Gladrow, Jannes
P23 (PH)	Prompt-Tuning SAM: From Generalist to Specialist with only 2048 Parameters and 16 Training Images	Piater, Tristan; Barz, Björn; Freytag, Alexander

12:30 End of conference