Energy Efficiency Analysis of Federated Learning: Insights from UAV-Based Thermal Imaging Applications

L. Duda[1], F. Alibabaei[1], V. Kozlov[1], L. Berberi[1], E. Vollmer[1], L. Klug[1], M. Benz[1], R. Volk[1], J.P. Gutiérrez Hermosillo Muriedas[1], M. Götz[1], J. Sáinz-Pardo Díaz[2], Á. López García[2], F. Schultmann[1], A. Streit[1]

Karlsruhe Institute of Technology (KIT),

Instituto de Física de Cantabria (IFCA)

***EGI2025**, 02-06. June 2025, Santander, Spain*

Funded by the European Union

Agenda

- Motivation & Foundations
  - Traditional machine learning
  - Thermal urban feature semantic segmentation
- Federated Learning
  - FL algorithms and approaches
- Frameworks
- Energy and accuracy results
- Conclusion

Traditional Machine Learning

- Train a model such that it recognizes a pattern or behavior
- Data is centralized in one spot
- “Data hunting” - the more data the better?
- GPT-3 (Generative Pre-trained Transformer 3) needed ≈ 570 GB of text data [1]

Challenges:

- What about distributed data that is not shareable?
- Unequally distributed data?

Thermal urban feature segmentation

- Identifying thermal anomalies (hot spots) in urban environments to improve the efficiency of energy-related systems [9]
- Images of Karlsruhe and Munich taken by drones at night

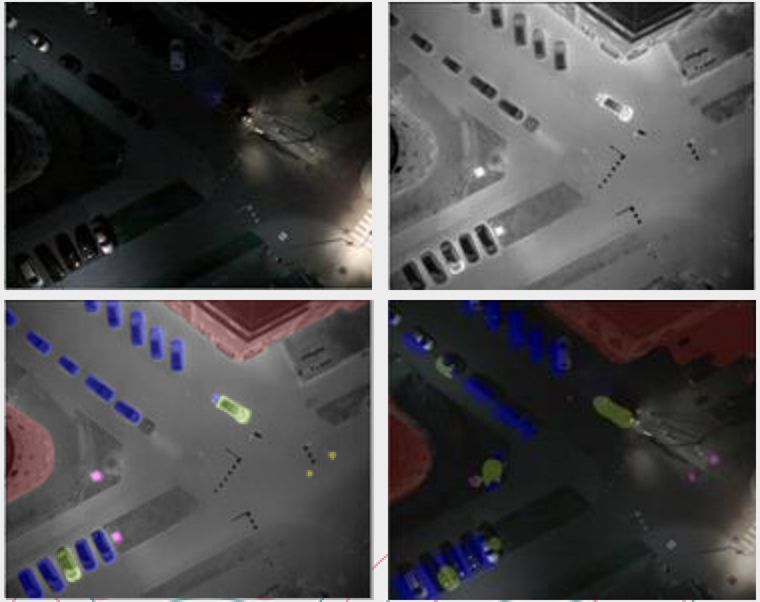

Example of thermal urban feature segmentation (I): combined RGB (top left) and TIR (top right) inputs, manual segmentation mask (bottom left), and U-Net model prediction (bottom right). Source [10]

Funded by the European Union

Thermal urban feature segmentation

- Semantic segmentation model
- U-Net [11] with ResNet-152 [12] backbone
- 700 images from Munich
- 93 images from Karlsruhe
- 8001 annotations

Class	# Annotations	# Pixel (*10^3)
Background	-	**37 063.96**
Building	**1404**	**9 087.95**
Car (cold)	**2531**	**601.90**
Car (warm)	**1034**	**325.60**
Manhole round	**1536**	**50.51**
Manhole square	**358**	**12.79**
Miscellaneous	**81**	**8.38**
Person	**275**	**7.64**
Street Lamp	**782**	**27.18**

Federated Learning

An approach enabling multiple peers to collaboratively learn a shared prediction model by sharing the weights of the model but not the data itself [2].

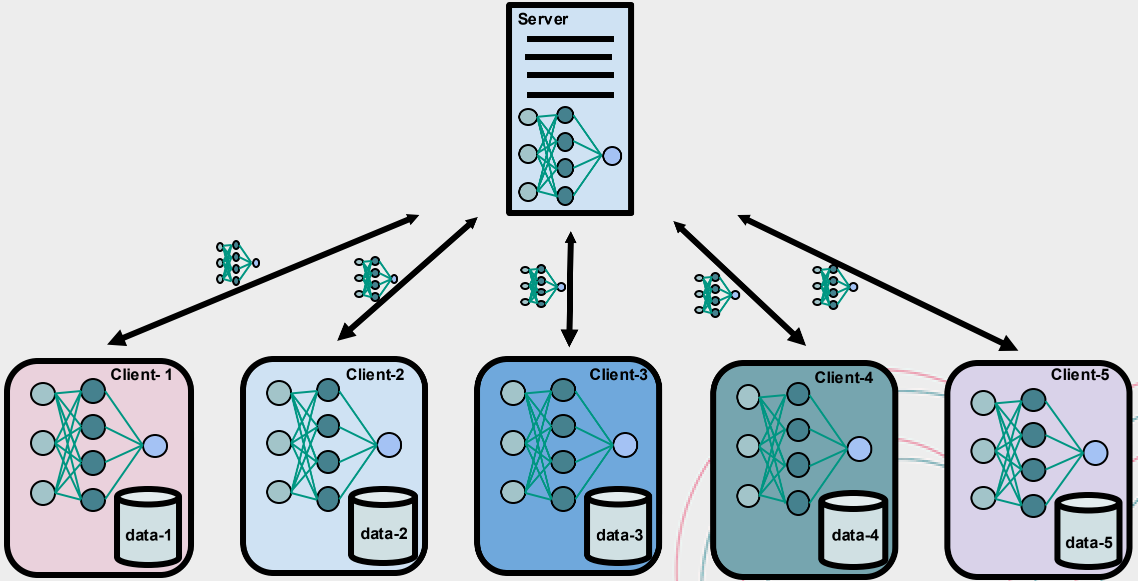

Scatter & Gather with different FL algorithms

- **FedAvg** [3]: Basic approach using weighted average for the aggregation of the updated weights
- **Fedprox** [4]: Extension of FedAvg that adds a regularization term to the local loss function to penalize the local weights that deviate from the global model
- **Fedopt** [5]: Added option of using a specified Optimizer on the client- and server-side when updating the model to improve the effectiveness
- **Scaffold** [6]: Added correction term to the model weights after each epoch of local training to prevent them from deviating too much from the global weights

Cyclic Weight Transfer (CWT) [7]

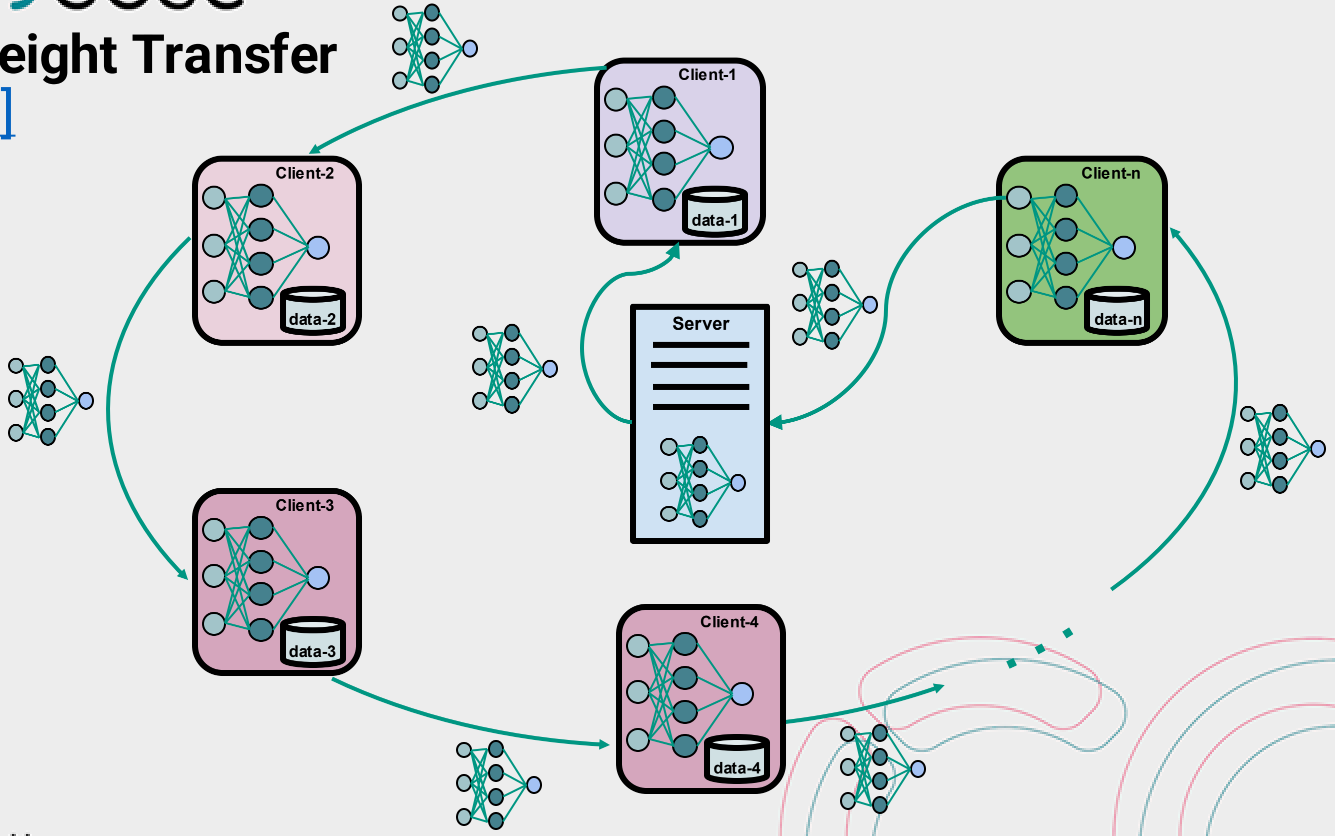

Swarm Learning (SL) [8]

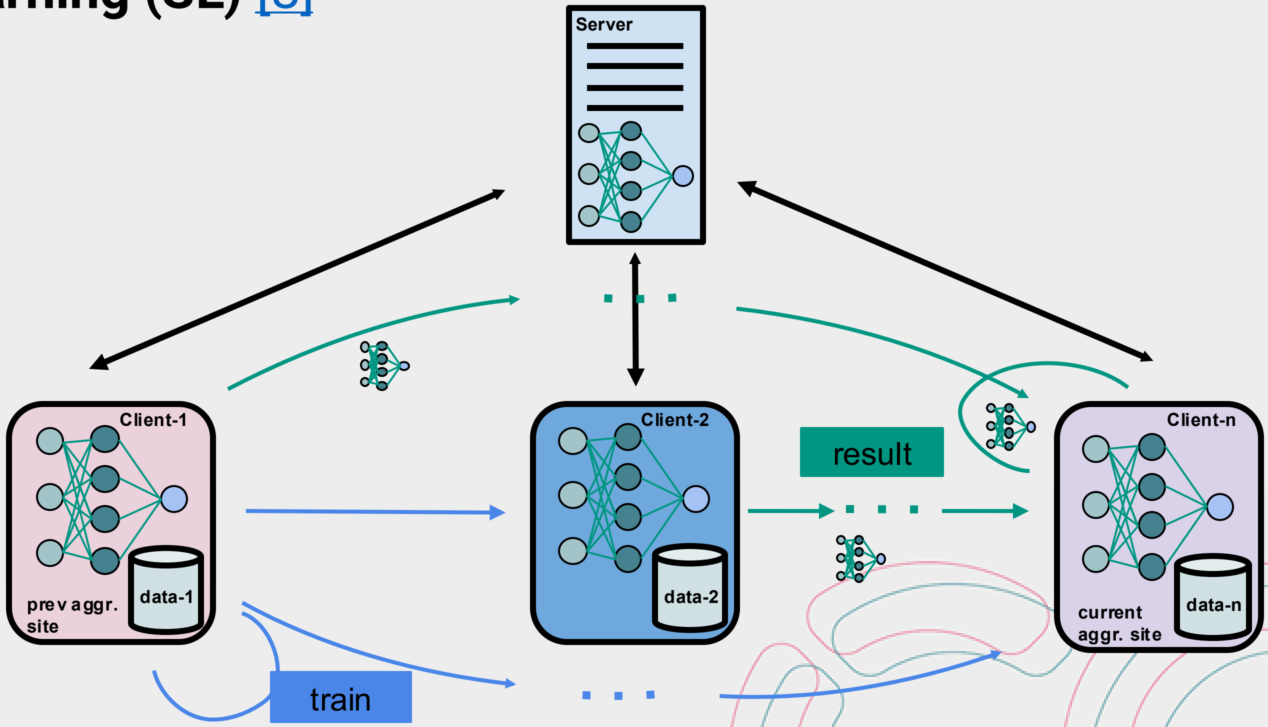

Frameworks

- NVFlare [13]: Open-source library
  - allowing to adapt existing machine learning workflows to a federated paradigm,
  - facilitating secure, privacy-preserving, and distributed multi-party collaboration.
  - Also available within the AI4EOSC Dashboard
- MLFlow [14] for experiment tracking, provided by AI4EOSC Project
- Perun [15] for calculating the energy consumption, provided by Helmholtz AI Energy

Accuracy metrics per algorithm

- CL shows the highest accuracy with FedAvg and Scaffold close
- FedOpt underperformed consistently across all metrics

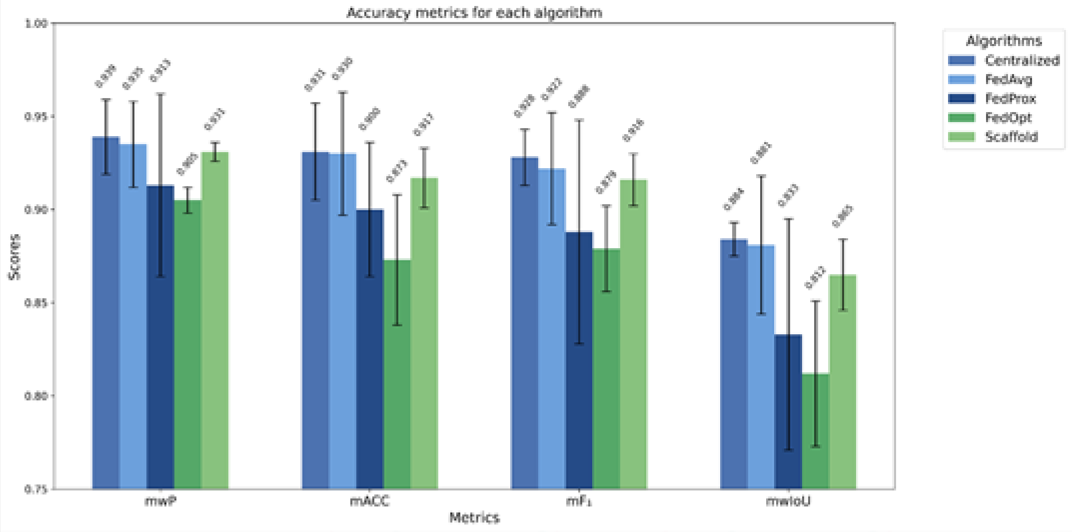

Energy metrics per algorithm

- FedAvg consumes ~164% more energy than CL FedOpt
- FedAvg takes ~162% more time to execute than CL
- Scaffold required the most training time and energy

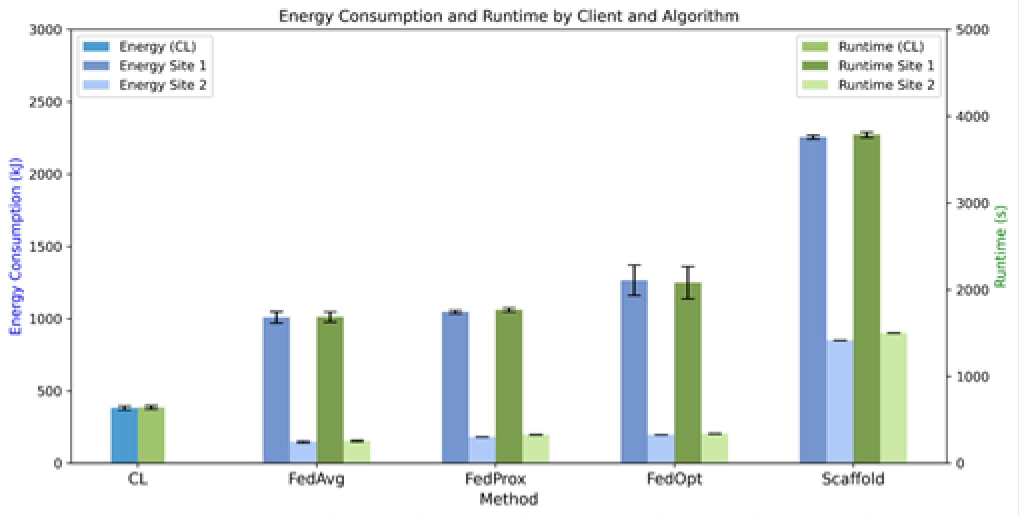

Accuracy metrics per workflow

- SL outperforms all other workflows including CL
- Client Order in CWT matters in regard of accuracy: Training KA (smaller dataset) first, then MU (larger dataset) improved generalization

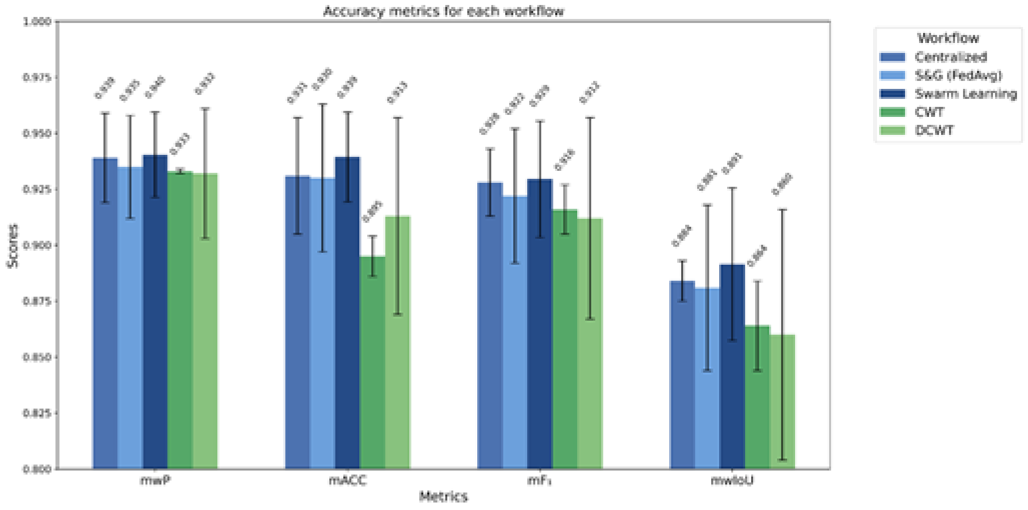

Energy metrics per workflow

- SL reduced FedAvg's training time by ~206% by eliminating server-client communication
- CWT required ~105% more time than SL due to sequential client training delays

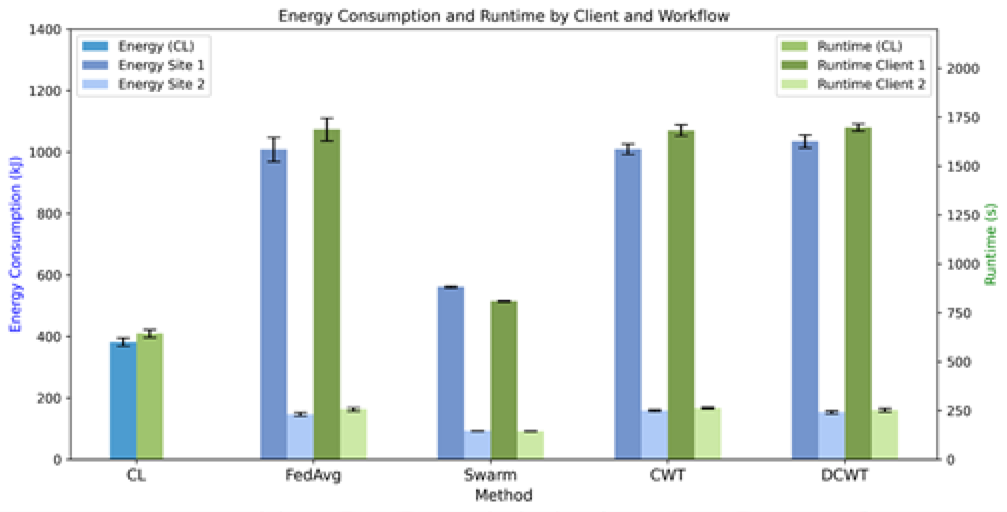

Funded by the European Union

Conclusion

- FL can consume more energy than Centralized Learning in some scenarios, but is also strongly depending on the algorithm and workflow
- Decentralized workflows (SL, DCWT) reduce training time
- Conversion of energy to carbon emissions depending on regional grid characteristics [16]

References

[1] Tom B. Brown, et al. (2020). Language Models are Few-Shot Learners. arXiv: https://arxiv.org/abs/2005.14165

[2] Roth, H. R., et al. (2022). NVIDIA FLARE: Federated Learning from Simulation to Real-World. *arXiv*. https://arxiv.org/abs/2210.13291

[3] H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, & Blaise Agüera y Arcas. (2016). Communication-Efficient Learning of Deep Networks from Decentralized Data. arXiv: https://arxiv.org/abs/1602.05629

[4] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, & Virginia Smith. (2018). Federated Optimization in Heterogeneous Networks. arXiv: https://arxiv.org/abs/1812.06127

[5] Sashank Reddi, Zachary Charles, et al.. (2020). Adaptive Federated Optimization. arXiv: https://arxiv.org/abs/2003.00295

AI4 eosc

References

[6] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank J. Reddi, Sebastian U. Stich, & Ananda Theertha Suresh. (2021). SCAFFOLD: Stochastic Controlled Averaging for Federated Learning. arXiv: https://arxiv.org/abs/1910.06378

[7] Chang, K., et al. (2018). Distributed deep learning networks among institutions for medical imaging. Journal of the American Medical Informatics Association : JAMIA, 25(8), 945–954. https://doi.org/10.1093/jamia/ocy017

[8] Stefanie Warnat-Herresthal, et al. (2021). Swarm Learning for decentralized and confidential clinical machine learning. https://doi.org/10.1038/s41586-021-03583-3

[9] Automated Thermography: https://ai4eosc.eu/use-cases/automated-thermography/

[10] Vollmer, E. (2023). UAV-based thermography: Using AI with multispectral data. Vortrag gehalten auf ANERIS Workshops on AI Basics for Image Processing (2023), Online, 28. November – 7. Dezember 2023. DOI: 10.5445/IR/1000166038

References

[11] Olaf Ronneberger, Philipp Fischer, & Thomas Brox. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation. arXiv: https://arxiv.org/abs/1505.04597

[12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, & Jian Sun. (2015). Deep Residual Learning for Image Recognition. arXiv: https://arxiv.org/abs/1512.03385

[13] NVFlare: https://github.com/NVIDIA

[14] MLflow: https://mlflow.org/

[15] Gutiérrez Hermosillo Muriedas, J. P., Flügel, K., Debus, C., Obermaier, H., Streit, A., & Götz, M. (2023). perun: Benchmarking Energy Consumption of High-Performance Computing Applications. In J. Cano, M. D. Dikaiakos, G. A. Papadopoulos, M. Pericàs, & R. Sakellariou (Eds.), Euro-Par 2023: Parallel Processing (pp. 17–31). Cham: Springer Nature Switzerland.

References

[16] Savazzi, S., Rampa, V., Kianoush, S., & Bennis, M. (2022). An energy and carbon footprint analysis of distributed and federated learning. IEEE Transactions on Green Communications and Networking, 7(1), 248–264.

EGI Conference 2025 - June 4th, 2025. Santander, Spain

Co-funded by the European Union

AI4EOSC

ai4eosc-po@listas.csic.es

ai4eosc.eu

Thank you for your attention!

Leonhard Duda - leonhard.duda@kit.edu
Fahimeh Alibabaei - khadijeh.alibabaei@kit.edu
Valentin Kozlov - valentin.kozlov@kit.edu
Lisana Berberi - lisana.berberi@kit.edu