



# SLOVENIAN AI FACTORY: SERVICES AND USE CASES



**Sašo Džeroski**

Head of Department of knowledge technologies

Jožef Stefan Institute, Ljubljana, Slovenia



# SLAIF BACKGROUND: KNOWLEDGE

## ***Slovenia has a strong AI community***

*AI research in Slovenia started in the mid 1970s*

- *Jozef Stefan Institute*
- *University of Ljubljana, Faculty of Computer Science*

## ***Slovenia has a strong HPC community***

- *SLING (Slovenian National Supercomputing Network)*
- *NCC (National Competence Center) as part of SLING*

***There is excellent collaboration between the two communities***  
*(Which is not to be taken for granted)*



# SLAIF BACKGROUND: INFRASTRUCTURE

## ***Extensive HPC infrastructure***

*Vega is one of the early EuroHPC systems. It connects Slovenia into Europe's high-end HPC ecosystem. Balanced architecture: Supports both CPU and GPU workloads. Equipped with datalake, excellent network connectivity to support data analytics and AI.*

## ***Established broader view of infrastructure (CLARIN, CLARIN.SI)***

*Data (e.g. language resources) are also infrastructure*

## ***Emerging view that AI is infrastructure (for science), even before the AI Factory concept was introduced (RAZUM@JSI)***

*In March 2024, we served a local instance of an LLM (Mistral) within JSI, both via a chat interface and API. We also support(ed) fine-tuning of LLMs with data curated by JSI colleagues.*

# SLAIF CONSORTIUM

*Research powerhouses, covering both HPC and AI*

- Jozef Stefan Institute
- University of Ljubljana



UNIVERZA  
V LJUBLJANI



*Infrastructure providers*

- IZUM (Institute of Information Science, Maribor)
- ARNES (Academic and Research Network of Slovenia)



*Knowledge institutions/Universities*

*(Maribor, Novo Mesto, Primorska, Nova Gorica)*



*Technology transfer organizations*

- Chamber of Commerce and Industry of Slovenia
- Technology Park Ljubljana



TEHNOLOŠKI PARK LJUBLJANA  
01

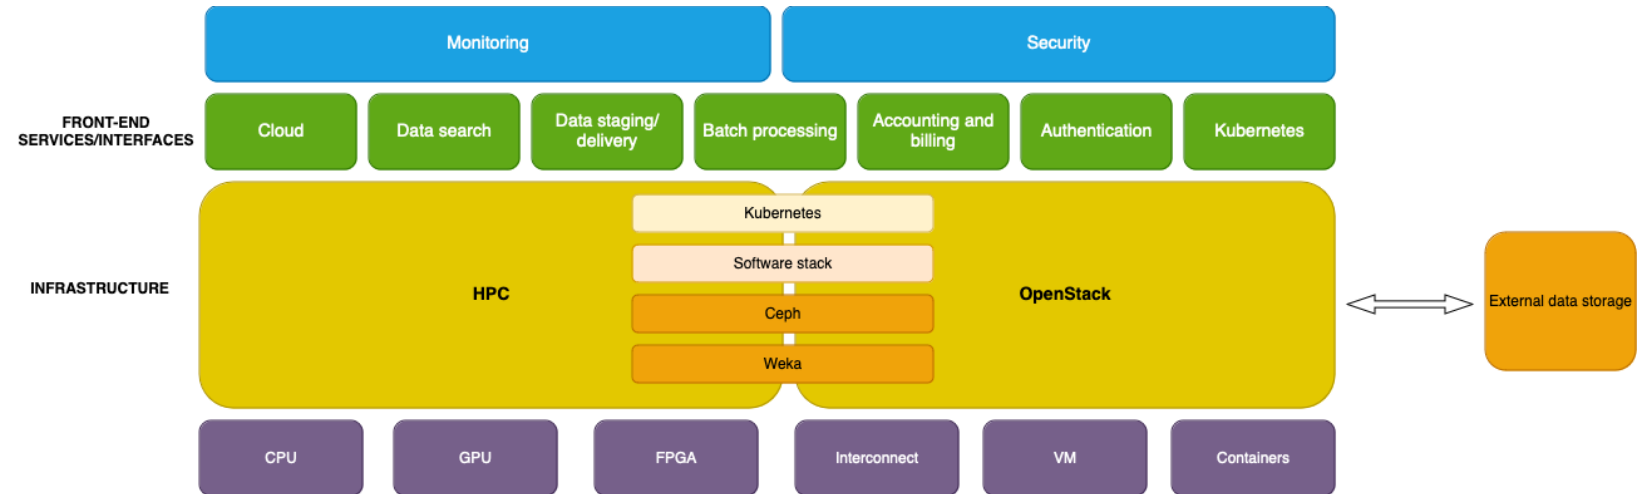
# SLAIF AI-HPC: AI-optimized Supercomputer



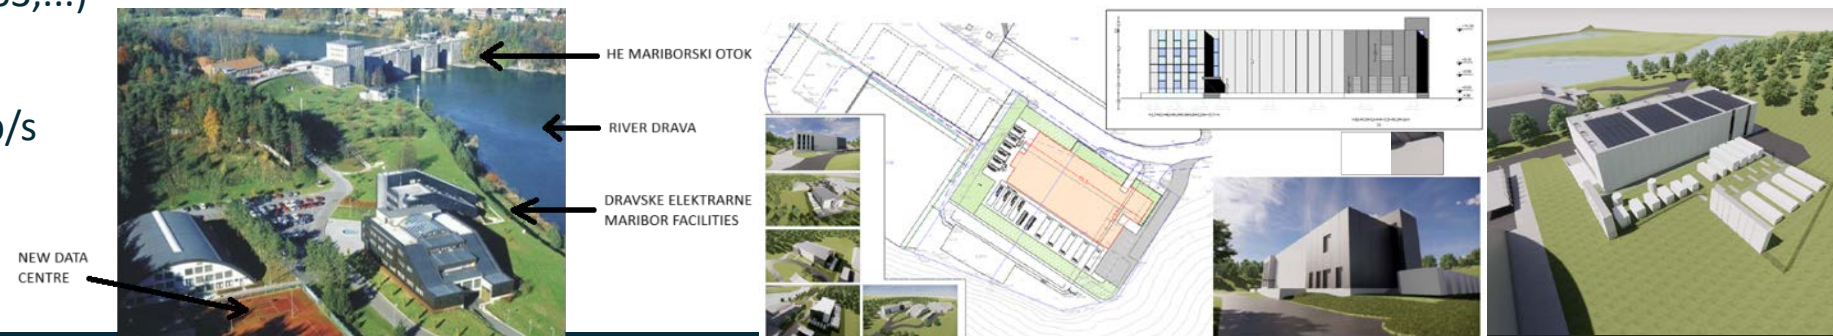
## Hardware:

- CPU
  - AMD (Venice) or
  - ARM64 (Grace, RHEA-2, Vera)
- GPU
  - NVidia GB200/300, Rubin
  - AMD MI450x
- Storage:
  - High throughput: Weka, VAST,
  - DDN Infinia...
  - Large capacity: CEPH (FS, S3,...)
- Interconnect:
  - Infiniband 800Gb/s, 1.6Gb/s
  - NVLink6/7 - 4/8/72

## Software:



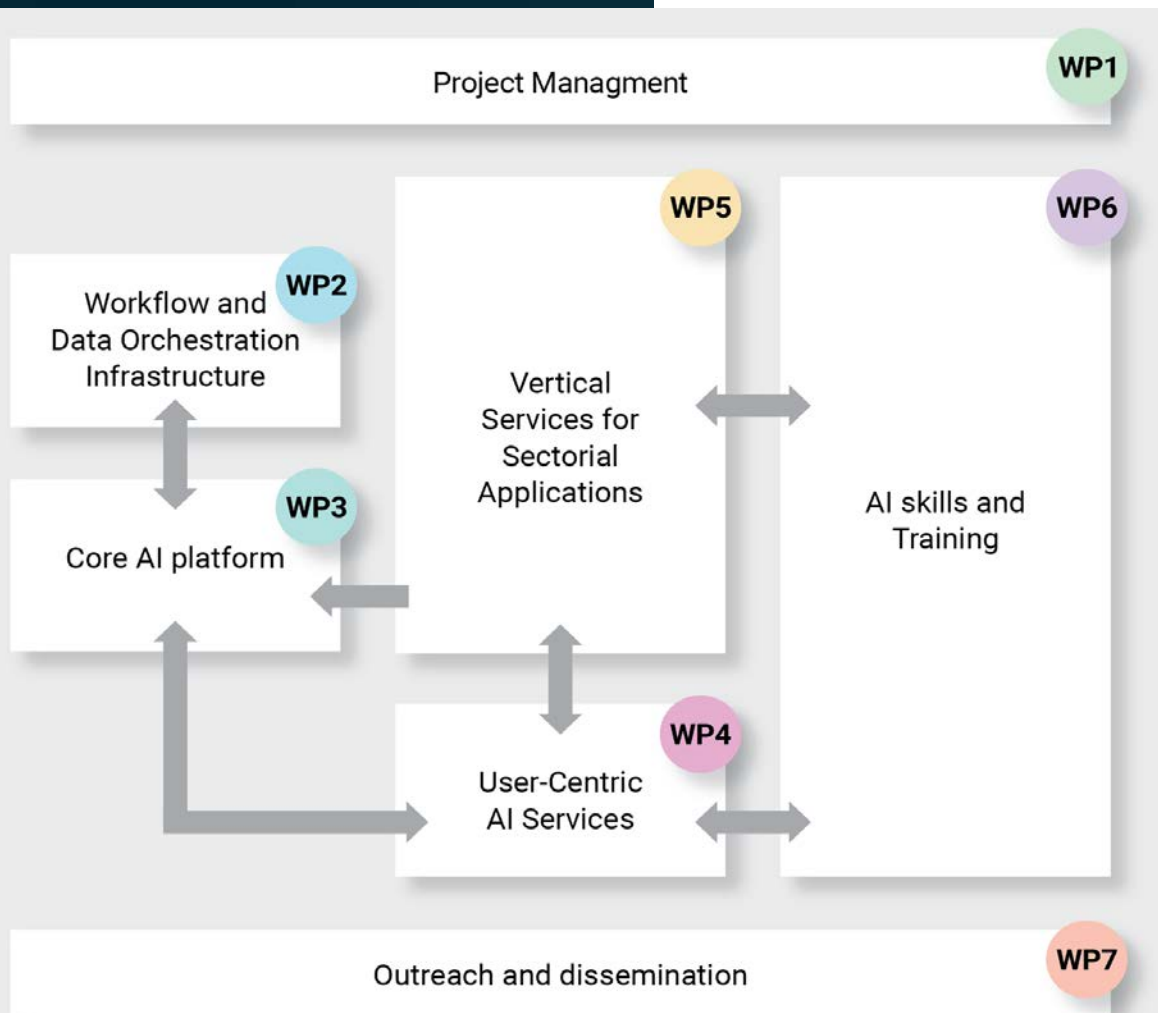
## Location:







# SLAIF AIF: The Slovenian AI Factory



## *Central WPs of the HE Project*

- *Workflow and Data Orchestration Infrastructure*
- *Core AI Platform*
- *User-centric AI Services*
- *Vertical Services for Sectorial Applications*
- *AI Skills and Training*



# SLAIF: CORE AI PLATFORM

## *Generic horizontal services*

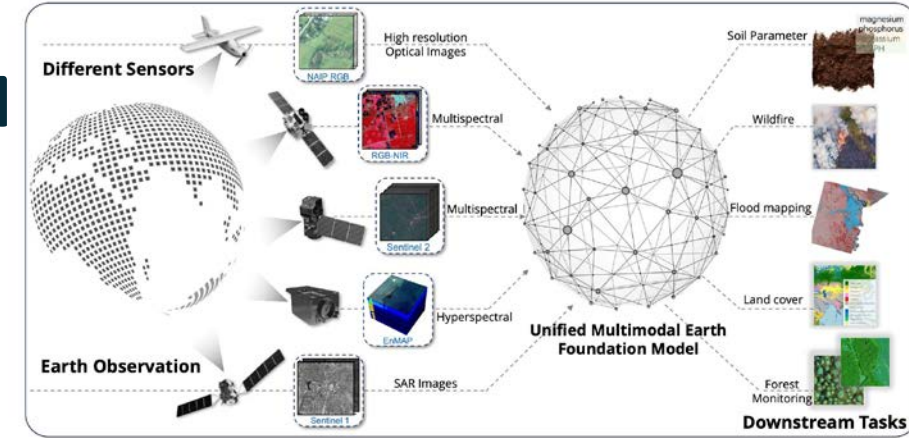
### *AI Workflows for*

- *Training, fine-tuning, evaluation, and inference for LLMs, VLMs, and multi-modal foundation models*
- *Explainable ML & AutoML*
- *Automated Scientific Modelling (of Dynamical Systems)*
- *AutoML/AutoOptimization/AutoAI*

### *Proof of Concept Solutions*

### *On-line Semantic Catalogue*

- *Making data/models/workflows/services FAIR*
- *Extending ontologies for describing data, ML algorithms, models*
- *Annotating artefacts, facilitating search & reuse*



# SLAIF: USER-CENTRIC AI SERVICES

## Data Storage Connectors

- *Establishing and maintaining data spaces*
- *Making AI-ready data available within the SLAIF platform*
- *An example will be remotely sensed data, starting w Slovenia*

## User Interface and Experience

*Interfaces for seamless connection to HPC/AI services*

*User-friendly access to workflows and the infrastructure*

## User Engagement and Feedback





# VERTICAL SERVICES FOR SECTORIAL APPLICATIONS

## Green Transition

- *Agriculture, Environment, Energy, Smart manufacturing*
- *Multi-modal foundation models for RSI relevant to A.&E.*

## Health & Biotech

- *Diagnostic, Personalized Medicine, Drug Discovery, ...*
- *Foundation models for sensor data (e.g. wearables) relevant*

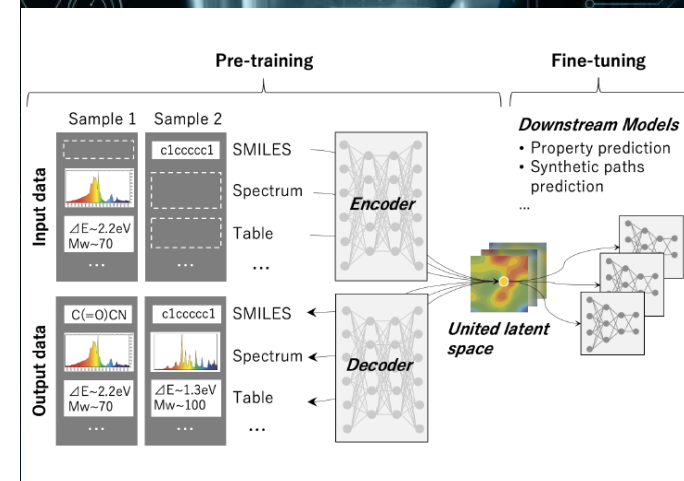
## Digital Society

- *Language technologies (LLMs for Slovene)*
- *GenAI for creative/media sector, Public admin, Education, ...*

## AI for Science

- *Life sciences, Materials sciences, Environmental sciences*
- *Interdisciplinary research*

**PoVeJMo**  
SLOVENSKI  
VELIKI JEZIKOVNI MODEL



# USE CASE: FINE-TUNING FOUNDATION MODELS FOR NUTRITION



## *The tasks:*

- ***Traffic Light Classification of Recipes:*** *healthy (green), moderate (orange), limit consumption (red)*
  - *Two datasets: Ingredients, Title+Ingredients*
  - *Four dimensions: Salt, Sugar, Fat, Saturates*
- ***Assessing Recipe Nutrient Values***
  - *Two datasets: Ingredients, Title+Ingredients*
- ***Making food data interoperable***
  - *Named entity linking (NEL)*
  - *Named entity recognition (NER) + NEL*
  - *Three datasets: CafeteriaFCD (recipes), CafeteriaSA (scientific abstracts), Artificial*
  - *Three ontologies: FoodOn, SnomedCT, Hansard*

# Fine-tune Llama3 with carefully curated datasets

*Zero-shot prompt:*  
[INST]Review the nutrient values per 100 grams in a recipe using these ingredients: 180 g wheat flours, bread, unenriched, 180 g wheat flour, white, cake, enriched, 200 ml sugars, granulated, 200 g butter, without salt, 200 ml water, bottled, generic, 1 pinch salt, table.[/INST]  
Nutritional values in each 100 g: energy - 364.03, fat - 17.96, protein - 4.09, salt - 0.05, saturates - 10.95, sugars - 18.33

*One-shot prompt:*  
[INST]The following are examples of questions (with answers) about nutrition. Question: Determine the nutritional profile per 100 grams in a recipe that uses these ingredients: 1 cup cheese, gouda, 4 tablespoon butter, without salt, 3/4 cup wheat flour, white, all-purpose, unenriched, 1/2 teaspoon salt, table, 1/2 teaspoon spices, pepper, red or cayenne, 1 tablespoon cream, fluid, heavy whipping Answer: Nutrient values highlighted for 100 grams: energy - 426.12, fat - 29.33, protein - 14.52, salt - 1.92, saturates - 18.33, sugars - 1.18 Respond to the following question in the same manner as seen in the examples above. Question: Review the nutrient values per 100 grams in a recipe using these ingredients: 180 g wheat flours, bread, unenriched, 180 g wheat flour, white, cake, enriched, 200 ml sugars, granulated, 200 g butter, without salt, 200 ml water, bottled, generic, 1 pinch salt, table Answer:[/INST]  
Nutritional values in each 100 g: energy - 364.03, fat - 17.96, protein - 4.09, salt - 0.05, saturates - 10.95, sugars - 18.33

*Five-shot prompt:*  
[INST]The following are examples of questions (with answers) about nutrition. Question: Gauge the nutrient values per 100 grams in a recipe prepared with the following ingredients: 2 cup cream, fluid, heavy whipping, 1 tablespoon spices, cardamom Answer: Per 100 g, the nutrient values are: energy - 339.02, fat - 35.36, protein - 3.04, salt - 0.07, saturates - 22.49, sugars - 2.85 Question: Establish the nutrient profile per 100 g in a recipe containing these ingredients: 1 tablespoon shallots, raw, 2 teaspoon spices, garlic powder, 12 cup peanut butter, smooth style, without salt, 3 tablespoon oil, sesame, salad or cooking, 2 tablespoon soy sauce made from soy (tamari), 1 teaspoon spices, ginger, ground, 1 teaspoon roland, seasoned rice wine vinegar, upc: 041224705142, 1/4-1/2 teaspoon spices, pepper, red or cayenne, 13 cup soup, chicken broth or bouillon, dry Answer: Nutrient profile for every 100 g: energy - 494.83, fat - 40.58, protein - 20.22, salt - 17.01, saturates - 8.28, sugars - 12.29 Question: Verify the nutrient values per 100 g in a recipe prepared with these ingredients: 16 ounce milk, fluid, 1% fat, without added vitamin a and vitamin d, 8 ounce beverages, almond milk, unsweetened, shelf stable, 13 cup sugars, granulated, 14 cup cornstarch, 12 teaspoon vanilla extract, 14 teaspoon shortening confectionery, coconut (hydrogenated) and or palm kernel (hydrogenated) Answer: Nutrient facts per 100 grams: energy - 340.40, fat - 1.30, protein - 0.41, salt - 0.03, saturates - 1.11, sugars - 50.82 Question: Identify the nutritional composition per 100 grams in a recipe with these ingredients: 500 g ground turkey, raw, 1 cup onions, raw, 12 cup bread crumbs, dry, grated, plain, 12 cup carrots, raw, 12 cup sauce, barbecue, 2 teaspoon sauce, worcestershire, 1 teaspoon spices, garlic powder, 34 teaspoon spices, pepper, black Answer: Nutrient profile for each 100 g: energy - 180.22, fat - 1.96, protein - 4.72, salt - 1.63, saturates - 0.42, sugars - 18.20 Question: Find the nutritional breakdown per 100 g in a recipe that uses the following ingredients: 12 pound pretzels, soft, unsalted, 21 ounce corn, sweet, white, raw, 12 ounce cookies, graham crackers, plain or honey, lowfat, 32 ounce nuts, walnuts, english, 7 ounce cookies, graham crackers, plain or honey, lowfat, 12 ounce cookies, graham crackers, plain or honey, lowfat, 1 pound butter, without salt, 12 drop sauce, ready-to-serve, pepper, tabasco, 8 tablespoon sugars, brown, 2 teaspoon spices, chili powder, 2 teaspoon sauce, worcestershire, 2-3 tablespoon spices, garlic powder Answer: The nutrient breakdown per 100 g is: energy - 383.98, fat - 14.10, protein - 7.92, salt - 0.59, saturates - 3.85, sugars - 3.92 Respond to the following question in the same manner as seen in the examples above. Question: Review the nutrient values per 100 grams in a recipe using these ingredients: 180 g wheat flours, bread, unenriched, 180 g wheat flour, white, cake, enriched, 200 ml sugars, granulated, 200 g butter, without salt, 200 ml water, bottled, generic, 1 pinch salt, table Answer:[/INST]  
Nutritional values in each 100 g: energy - 364.03, fat - 17.96, protein - 4.09, salt - 0.05, saturates - 10.95, sugars - 18.33

[INST] Verify the FSA traffic light labels per 100 grams in a recipe that has these ingredients: 1 cup water, bottled, generic, 1 tablespoon oil, olive, salad or cooking, 34 teaspoon salt, table, 3 cup wheat flours, bread, unenriched, 2 teaspoon sugars, granulated, 1 teaspoon wheat flours, bread, unenriched, 2 teaspoon spices, parsley, dried, 1 teaspoon semolina, unenriched [/INST]  
Food Standards Agency traffic lights per 100 g presented: fat - green, salt - red, saturates - green, sugars - green

[INST] Evaluate the FSA traffic light labels per 100 grams in Easy Pizza Dough for Bread Machine using these ingredients: 1 cup water, bottled, generic, 1 tablespoon oil, olive, salad or cooking, 34 teaspoon salt, table, 3 cup wheat flours, bread, unenriched, 2 teaspoon sugars, granulated, 1 teaspoon wheat flours, bread, unenriched, 2 teaspoon spices, parsley, dried, 1 teaspoon semolina, unenriched [/INST]  
The fsa traffic lights per 100 g listed are: fat - green, salt - red, saturates - green, sugars - green

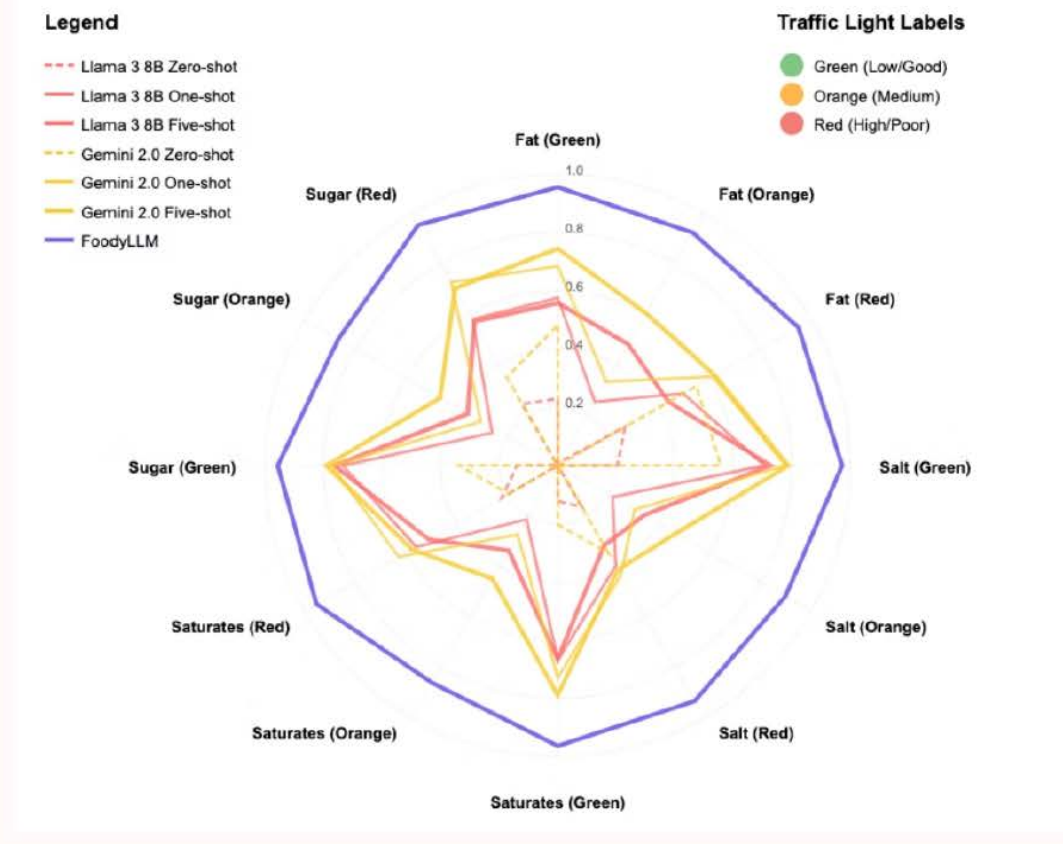
## EXAMPLE PROMPTS FOR THE THREE TASKS

[INST] Find all food entities present in the text: Mix the cream cheese, beef, olives, onion, and Worcestershire sauce together in a bowl until evenly blended. Keeping the mixture in the bowl, scrape it into a semi-ball shape. Cover and refrigerate until firm, at least 2 hours. Place a large sheet of waxed paper on a flat surface. Sprinkle with walnuts. Roll the cheese ball in the walnuts until completely covered. Transfer the cheese ball to a serving plate, or rewrap with waxed paper and refrigerate until needed. [/INST]  
Certainly, the entities connected with food are outlined as follows: cream cheese, beef, olives, onion, Worcestershire sauce, walnuts, cheese ball.  
[INST] Is it within your capabilities to link the extracted food entities to a Hansard taxonomy? [/INST]  
Certainly, the entities are associated properly: cream cheese - AG.01.e [Dairy produce]; AG.01.e.02 [Cheese]; AG.01.n [Dishes and prepared food]; AG.01.n.18 [Preserve], beef - AG.01.d.03 [Beef], olives - AG.01.h.01.e [Fruit containing stone], onion - AG.01.h.02.e [Onion/leek/garlic], Worcestershire sauce - AG.01.h [Fruit and vegetables]; AG.01.l.04 [Sauce/dressing], walnuts - AG.01.h.01.f [Nut], cheese ball - AG.01.e.02 [Cheese]; AG.01.n.18 [Preserve].  
[INST] Might I trouble you to connect the extracted food entities to a FoodOn ontology, if possible? [/INST]  
Definitely, the entities are linked suitably: cream cheese - FOODON-03301889; FOODON-00001013, onion - FOODON-03301704; NCBITaxon-4679, worcestershire sauce - FOODON-03305003; FOODON-03311146, walnuts - NCBITaxon-16718, cheese - FOODON-00001013.  
[INST] Might it be possible for you to link them to a SNOMEDCT ontology, if needed? [/INST]  
Absolutely, the entities are related as such: cream cheese - SNOMEDCT-226849005; SNOMEDCT-255621006; SNOMEDCT-102264005, beef - SNOMEDCT-226916002, olives - SNOMEDCT-227436000, onion - SNOMEDCT-735047000, worcestershire sauce - SNOMEDCT-443701000124100; SNOMEDCT-227519005, cheese - SNOMEDCT-102264005.

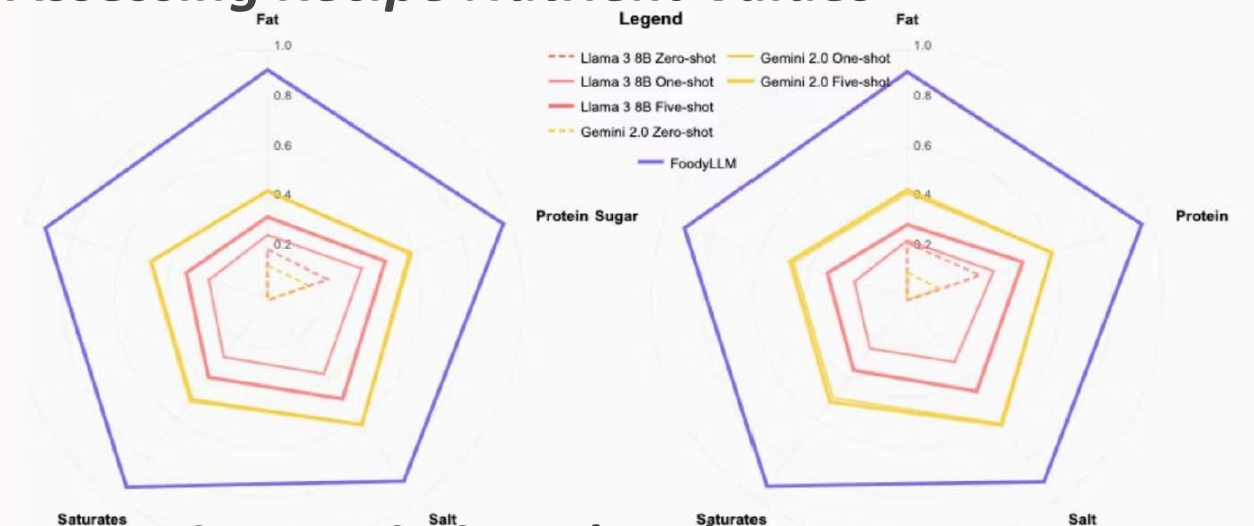


# FoodyLLM Results

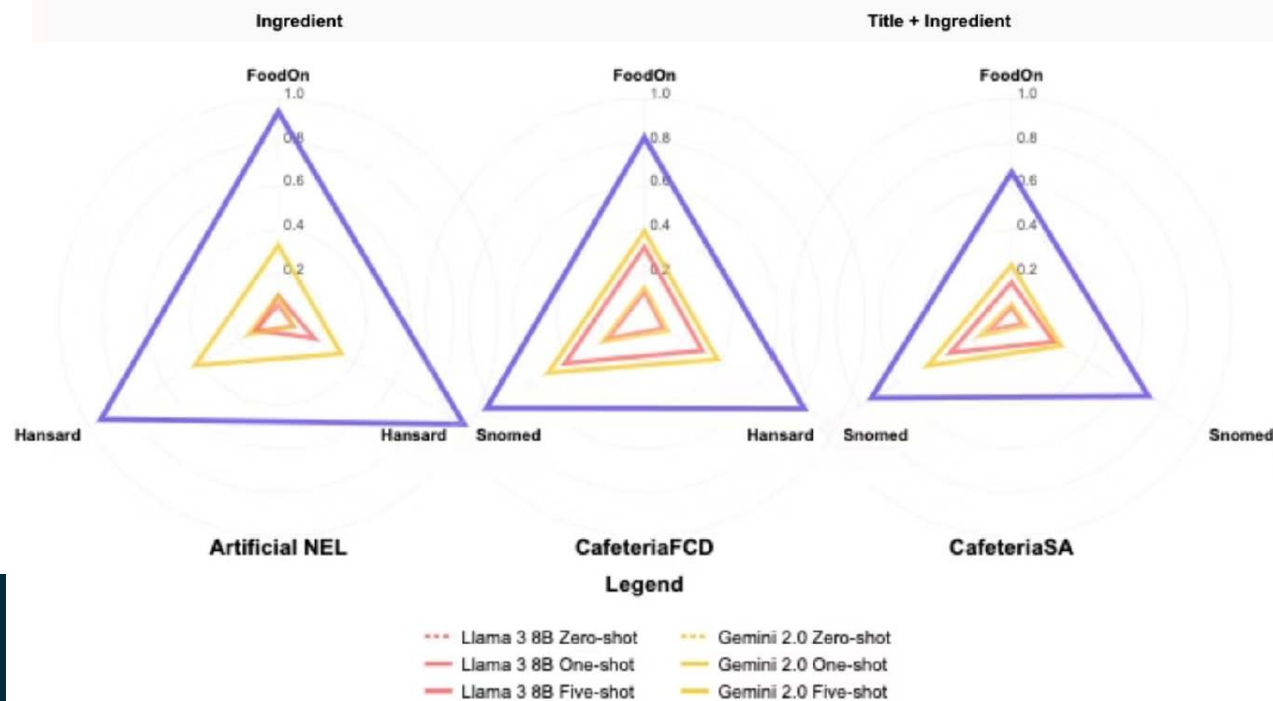
## 1. Traffic Light Classification



## 2. Assessing Recipe Nutrient Values



## 3. Named entity linking (NEL)



# MACHINE LEARNING FOR MATERIALS DESIGN

## *Data*

*Parameters of the material synthesis process*

- Temperature, Precursor concentration, Medium (acidity, ligands), Time, ...

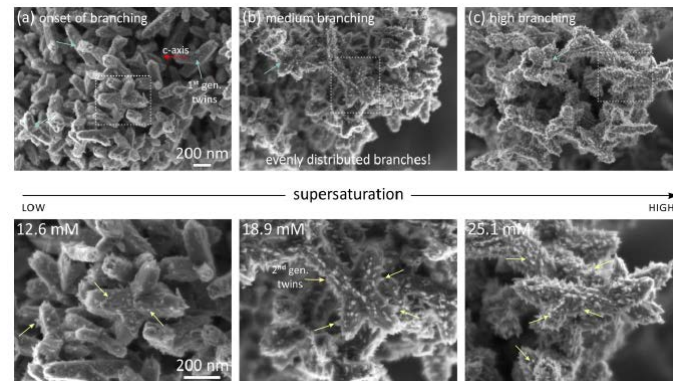
*Quantum-level properties*

- DFT simulation results
- Band-gap

*Morphology of the material (images)*

*Functional properties of materials*

- Electrical conductivity
- Catalytic activity
- Magnetic properties
- Anti-corrosive properties



## *Machine learning tasks*

*Find relations between the different kinds of data*

Primarily predict functional properties of materials from the other kinds of data (synth. parameters, morphology, quantum-level properties)

But also find relations between other pairs/ combinations of data, e.g.,

- *Parameters of synthesis and*
- *Morphology of the material*
- or
- *Morphology of the materials and*
- *Properties of the material*



# COMBINING ML & OPTIMIZATION FOR MATERIALS DESIGN



## *The approach*

Use **machine learning** to learn a model **M** predicting material properties (**P**), from synthesis parameters (**S**):  
 $P = M(S)$

The task of material design is to find a set of synthesis parameters' values **s** that optimize the values of the properties **P**

## *The outcome*

Suggested sets of synthesis parameters' values that can be used to synthesize new materials

For this task, we can use an **optimization** algorithm over the space of synthesis parameter values

- If we are optimizing the value of one property, single-criterion optimization; For multiple properties, we need to use multi-criteria/objective optimization (MOO)

## *Example application: Foamed glasses*

Predict five properties of the foamed glass, based on the eight parameters of the synthesis process

Process parameter	Material property
Water glass content [wt.%]	Apparent density $\rho_{app}$
Carbon black content [wt.%]	Pycnometric density $\rho_{pyc}$
Mn <sub>3</sub> O <sub>4</sub> content [wt.%]	Overall porosity $\epsilon_{total}$
Furnace temperature [°C]	Closed porosity $\epsilon_{closed}$
Foaming time [min]	Open porosity $\epsilon_{open}$
K <sub>3</sub> PO <sub>4</sub> content [wt.%]	
Drying	
Mixing	

A multi-target prediction tasks, where we have used semi-supervised and active learning

New materials synthesized based on suggestions



# DEVELOPING THE (SLOVENIAN) AI ECOSYSTEM

## *AI for Science*

*Identified as area of strategic interest for the EU. European Strategy for AI in Science (RAISE) released last week (OCT 8). Launch at policy event **AI in Science Summit**, Copenhagen (NOV 3-4): [ais25.eu](https://ais25.eu)*

*Also organized AI for Science scientific event ([ai4science.si](https://ai4science.si)), SEP 22-26, Ljubljana, including AI factories track*

## *AI for Society (incl. industry and public sector)*

- Business conference of Slovenian automotive industry (OCT 10)*
- Challenges and opportunities for Slovenian robotics (OCT 16)*
- Coming up: SLAIF Open Industry Day (NOV 17)*

*With the goal to better understand industry needs and appropriately plan the development of the Slovenian AI ecosystem*