Deliverable 5.1: Plan & report of AI weed mapping for AWM (version 1)

Grant Agreement number: 101083589
Start date of the project: 01/05/2023 **End date of the project**: 30/04/2027
Deliverable due date: 30/04/2024 **Date of delivery**: 30/04/2024
Classification: Public
Associated Work Package(s)

WP1	WP2	WP3	WP4	WP5	WP6	WP7	WP8
☐	☐	☐	☐	☑	☐	☐	☐

Version History

Version number	Implemented by	Notes
1.0	UG	

Table of contents

Funded by the European Union

Funded by the European Union under Grant Agreement No. 101083589. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or REA. Neither the European Union nor the granting authority can be held responsible for them.

Image processing activities and AI pipeline for weed detection

Work package (WP) 5 of the GOOD project deals with digitalization technologies of agroecological weed management systems. Here, we report on the progress on Task 5.1, automated weed identification and mapping. This automated weed mapping can eventually be applied for site-specific weed management. Very high (mm) to high (cm) resolution RGB data are collected with UAS (unmanned aerial system) technology and are then analysed using deep learning technologies, either in a semi-automated (supervised) or automated (unsupervised) way. This results in weed density maps of either all weeds combined, of different genera of weeds, or of different species of weeds. The latter aspect depends on the exact requirements for weed control in the individual Living Labs, which are introduced below. The application is also different in annual and perennial crops.

Introduction: experimental design

Annual crops

For each LL's main crop, three cover crops and one control plot (without cover crop) will be established in the first year. This 1st year is a pre-selection stage which will allow to select the most suitable cover crop for further experiments in year 2 and 3, with or without AMF inoculation.

The use of cover crops will be combined with other weed control strategies applied to the main crop. For each pilot site, these treatments will be based on the levers described in Section 2, and include, at least, 1) one cultural, 2) one mechanical, 3) one chemical "standard chemical practice"), 4) one chemical at reduced rate, and 5) one weedy treatment ("no-weeding control").

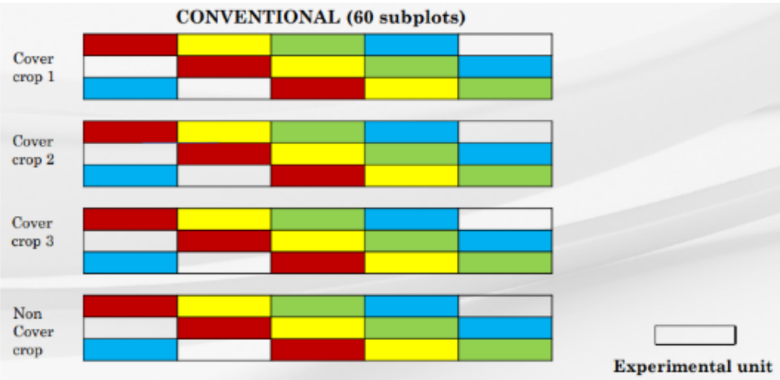

Perennial crops

Three cover crops (single species or mixtures: legume, grass, cereals, etc.) are sown on the interrows (corridors) of the orchard or vineyard; a reference is kept without cover crop ("interrows standard practice"). Four additional practices are experimented on the interrows: herbicide–full rate; herbicide–half rate; one nonchemical practice (e.g., mulching, mowing, etc.), plus an untreated control (weedy).

There will be 7 treatments in total for the conventional pilot sites with a minimum of 3 replications per treatment (21 subplots).

The treatments for both the conventional and organic & mixed sites are specified in the Living labs calendar (available on SharePoint).

Important weeds and level of discrimination detail of LLs

A questionnaire was sent to each LL to inquire about the exact purpose of the weed mapping system; feedback was received from all but two LLs. The outcome per LL is presented in **Table 1**. This reveals that for an efficient weed management, in all LLs, weeds should be discriminated to at least the genus level. This level of discrimination determines to a large extent the required resolution of the UAS data,

Funded by the European Union

Funded by the European Union under Grant Agreement No. 101083589. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or REA. Neither the European Union nor the granting authority can be held responsible for them.

since it is not only sufficient to detect weeds as such, e.g. as plants (green objects) in between rows, but for species/genus level discrimination, the resolution should be sufficient to distinguish between genera, looking at the shapes of the leaves. In row crops, where weed treatment is typically done early after germination, this requires a resolution of at least 5mm.

Further, the enquire also confirmed the diversity of the weed species to be discriminated within the project. As expected, there are large differences between the species that should be discriminated according to region and crop type.

Table 1: results questionnaire of the most damaging weeds and required level of distinction for each living lab

Country	Crop	Code	Most important weed species	Classification detail
Cyprus 1	Olives	CY_olives	/	/
France	Apple-Plum	FR_apple-plum	Broadleaf weeds, monocotyledon weeds	Species, at least dicotyledons vs monocotyledons
Greece1	Wheat	GR_wheat	*Avena, Lolium, Sinapis* and *Galium*	Species
Greece2	Grapes	GR_grapes	All the perennial weeds	Species
Italy1	Triticale	IT_triticale	*Chrysanthemum coronarium*, *Oxalis cernua*, *Malva* sp.	Genus
Italy2	Citrus	IT_citrus	*Oxalis pes-caprae* L. - *Urtica dioica* L.	Species
Italy3	Grapes	IT_grapes	*Cynodon dactylon* L.	Species
Latvia	Rye-Pea	LV_rye-pea	*Agropyron repens*, *Galium aparine*, *Matricara perforate*, *Poa annua; Bromus secalinus*	Genus
Netherlands	Onion	NL_onion	*Chenopodium album, Poa annua, Galinsoga* spp.	Species
Portugal1	Cowpea	PT_cowpea	*Cytisus scoparius, Cistus ladanifer, Lavandula stoechas*	Species
Portugal2	Olives	PT_olives	*Cytisus scoparius, Cistus ladanifer; Lavandula stoechas*	Species
Serbia1	Soybean	RS_soybean	*Chenopodium album, Sorghum halepense*	Species
Serbia2	Maize	RS_maize	*Chenopodium album, Sorghum halepense*	Species
Spain1	Rice	ES_rice	*Echinochloa* spp., *Leptochloa* spp, *Cyperus*, *Ammania*	Species
Spain2	Cherry	ES_cherry	*Lolium*, *Echium*, *Sonchus*, rumex	Genus
Spain3	Apple-Grapes	ES_apple-grapes	/	/

Data collection

The image gathering is realised through the use of UAS (unmanned aerial systems) and using high resolution RGB data. Practical guidelines for performing the flights were compiled and were shared

Funded by the European Union

Funded by the European Union under Grant Agreement No. 101083589. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or REA. Neither the European Union nor the granting authority can be held responsible for them.

with all LLs in a document. The separate LLs were also contacted for further guided support with the organisation of the UAS flights. The goal of the document was to ensure that image collection was done in a standardized way, and that all data were of high quality and readily available for post-processing. This guideline includes the following information:

Flight approach: Mapping flights

It was decided (see further) that we want to work with orthomosaic imagery, a single composite of the individual images covering the entire field, and not with individual UAS images. This requires for a mapping flights, with sufficient overlap to process the individual imagery and convert them into orthomosaics. Mapping flights are best executed automatically making use of the waypoint controls software. This can be done by a preprogrammed flight using the DJI Pilot, DJI Pilot 2 or other dedicated software. It is important to use sufficient overlap between the images, with at least 70% in horizontal and 70% in vertical direction; 75% if feasible.

Camera resolution

To ensure the ability to map the weeds, a minimal resolution of 20 MP is needed. For this application, possible drone-sensor combinations are the DJI Mavic 2 series DJI Mavic 3 series (DJI Mavic 3 Pro; DJI Mavic 3 MS), or Zenmuse P1 gimbal on DJI M300 or DJI M350; other combinations are of course also possible. Note that the resolution is also determined by the flight height (see further), and that the quality of the imagery is not only a function of the resolution of the sensor, but also of its intrinsic qualities (sensor size; optical quality, data storage, sensor settings, flight conditions, etc.)

Flight height

The discrimination between weeds depend on high resolution images. To obtain this, the minimal height possible for a flight should be selected. For a standard mapping flight with DJI software, this is 12m above the ground.

It is important to note that safety always should be the main priority. At low flight heights, special attention need to be paid to the presence of obstacles (trees, poles, buildings) surrounding the field. If required, either the flight height must be increased to safe levels (at least 3 m above objects that need to be overflown) or the area needs to be adjusted (keeping a safe distance of at least 5m from taller objects).

Ground control points

Ground control points (GCPs) are reference panels put in the field to extract the exact location for the orthophoto creation. Commonly, a 2x2 black-and-white checkerboards are used, of which the location of the centre is extracted with a RTK-GNSS system. It is recommended to put 6 to 8 GCP's per flight, evenly spread out over the flight area. The use of GCPs is important

Flight conditions

If possible, fly close to solar noon (less shadows) or in overcast (but still bright enough) conditions. Reduce the flight speed if conditions are cloudy, to avoid motion blur. Ideally, keep the shutter speeds small (<1/1000 of a second) and your ISO setting also relatively low. In changing weather conditions, don't use manual, or set your camera to AUTO-ISO, so that every image is bright enough. Set your camera to -0.5 or -0.7 stops to avoid saturation of brighter plants.

Data are collected by each individual LL following these guidelines.

Image processing

The collected data consist of a set of individual images per executed flight. In GOOD, we decided to work with the orthomosaic rather than the individual images, although this requires an additional processing stage. The reason for this choice is that the orthomosaic option has several advantages: *i*)

Funded by the European Union

Funded by the European Union under Grant Agreement No. 101083589. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or REA. Neither the European Union nor the granting authority can be held responsible for them.

every single individual plant is imaged in a near-nadir position, offering a more homogeneous dataset; *ii*) the total size of the dataset is much reduced, which is more efficient for data labelling, handling and storage; *iii*) the orthomosaic is precisely georeferenced, and therefore all products generated with the mosaic will also be accurately georeferenced, and *iv*) this orthomosaic can be processed immediately to field maps of weed density and further on to weed prescription maps, as was promised in the GOOD project proposal. Apart from site-specific weed management, georeferenced maps also allow precise monitoring of weed populations through time. It was decided that UG will generate the georeferenced orthomosaics for all LLs. Starting from this orthophoto, several processing steps will be performed, to prepare smaller, individual images, suitable for modelling.

Orthomosaic generation

To make the process clear, we will use a specific example (Fig 1a). Using a DJI Mavic 2, a typical flight can give rise to about 400 images, depending on the height and area of the flight. Each image with this drone is 20MP.

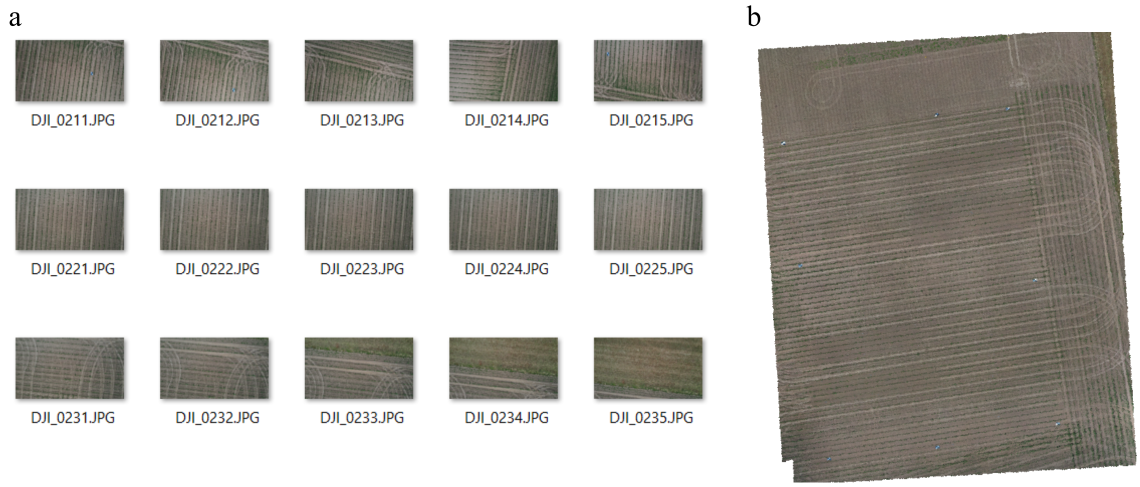

Figure 1: a) Example of several images taken by the DJI Mavic 2 drone; b) the resulting orthomosaic, covering an area of 76 by 80m.

Image mosaicking

The individual images are converted into the orthomosaic and digital elevation model (DEM) using dedicated structure-from-motion software, in this case Agisoft Metashape Professional (Version 1.8.4) will be used. All images have a certain degree of overlap. Correspondence between the images is found using feature detection, commonly scale invariant feature transform (SIFT). Once the features have been detected, they are matched between different images. The first step of the process is the image alignment, resulting in a sparse 3D point cloud and a first estimate of the camera positions. After this, GCPs can be used to georeference this alignment precisely and further correct the camera positions. The next steps include the generation of a dense 3D point cloud, mesh, digital elevation model, and finally a georeferenced orthomosaic that can be exported (Fig. 1b). Unless processed otherwise, this orthomosaic will have the same ground sampling distance (ground resolution distance) as the original images.

Funded by the European Union

Funded by the European Union under Grant Agreement No. 101083589. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or REA. Neither the European Union nor the granting authority can be held responsible for them.

Image annotation

Image annotation (labelling) is important for two reasons. First, it allows developing a supervised deep learning algorithm using these labels. Second, it also allows to evaluate the performance of such models, as well as of semi-supervised or unsupervised models. Manual labelling is labour-intensive; therefore, only a small subset of the actual orthomosaic can be labelled. The orthomosaic is cut up in smaller pieces (tiles) to allow smooth image annotation. These tiles will be uploaded on CVAT (a user interface set up by Eden) for annotation and then the annotated images will be available on Eden Library website.

Guidelines

The guidelines for the LLs on the annotation are provided by Eden and are described in Instructions for annotation with CVAT.pdf (available on SharePoint under Task 5.1). Once the orthophoto of a specific flight is generated, the optimal annotation strategy will be discussed with the LL. This strategy depends on a few factors such as:

1) The resolution of the generated orthophoto
2) The estimated amount/density of weeds in the field
3) The specific configuration of the plots, involved in the pilot field

The annotation happens with the free web-based tool Computer Vision Annotation Tool (CVAT). This tool allows to annotate specific regions that the user defines as objects. In Figure 2, an example is given of this process, where the defined objects are weed plants in between a main crop (maize). For this project, bounding boxes are used to define the location of the weeds. It is very important that every weed in an image is annotated, as its performance heavily depends on abundant and correct annotations.

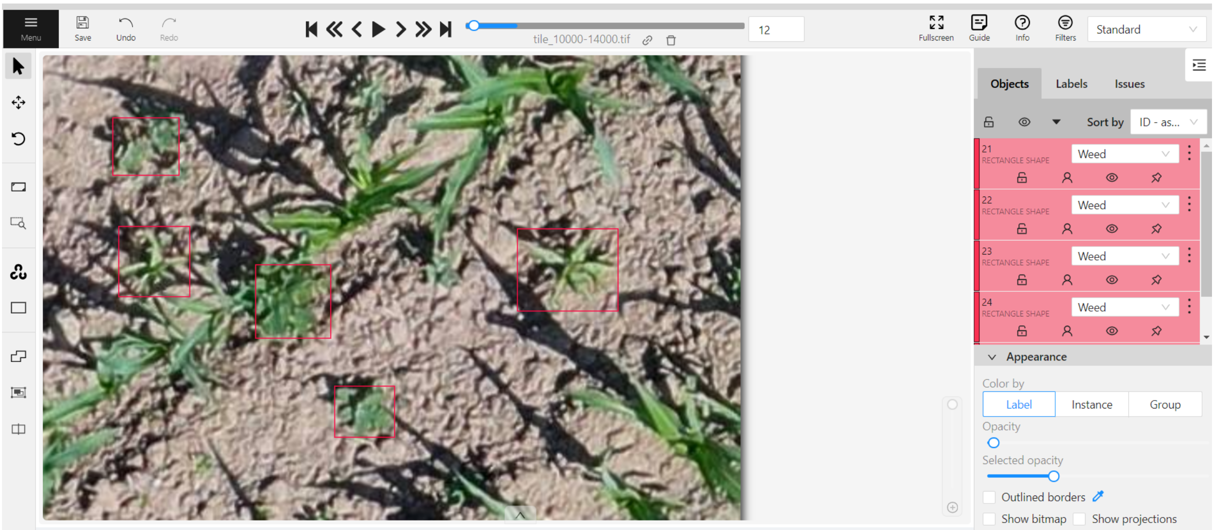

Figure 2: example of the annotation process for a maize dataset.

Sampling strategy

UG and Eden will collaborate to present the sampling strategy and sampling input data (i.e., to select the tiles that need to be sampled). This subset will be based on two general ideas:

1) The experimental design of the specific pilot field
2) Input from the LLs weed scientists, with regard to the most important weed patches to annotate.

The subsampling needs to take the experimental units into account. With this information, a stratified random sampling will be used to make sure that every experimental unit has at least one, and preferably several, annotated tiles. Also, the weed scientists will give input to which experimental units are more useful so sample annotated tiles from. As an example, experimental units without treatments might show a more diverse and more interesting weed flora in comparison to the application of full herbicide rate.

Funded by the European Union

Funded by the European Union under Grant Agreement No. 101083589. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or REA. Neither the European Union nor the granting authority can be held responsible for them.

Each tile will consist of a predefined size, probably covering about 40x40 cm. In this tile, all weed plants will need to annotated.

Annotation format

The annotations, made with the CVAT tool, will be exported in a standardized format. This is set to be the YOLO format. This format actually consists of a zip file, that holds the location of the annotated bounding boxes as a text file, for each annotated image.

Modelling

The goal is to map the spatial distribution of the weeds present in the pilot field. This includes the detection, as well as the identification.

Input data

The dataset defined by the stratified random sampling will be used as input for a deep learning model. This consists of the tiles, each with their annotation file.

Supervised model workflow

The general workflow for a supervised neural network is given in Figure . The selected tiles through the stratified sampling (dataset) are split into three sets, training data, validation data and test data. The training data is used as input for the model. This model is updated by stochastic gradient descent (SGD), based on the loss function. Parallel to that, validation data is used to evaluate the model, without using this data to update the model. Finally, test data is evaluated by the model after the training finished. This workflow guarantees that a model does not overfit on the given training data. When the performance of the model is favourable on the test data, the model can be used to provide a weed mapping of the non-annotated tiles. Further, the goal is to develop a model that can also be applied to future datasets of this crop and region (without requiring new training or new annotations).

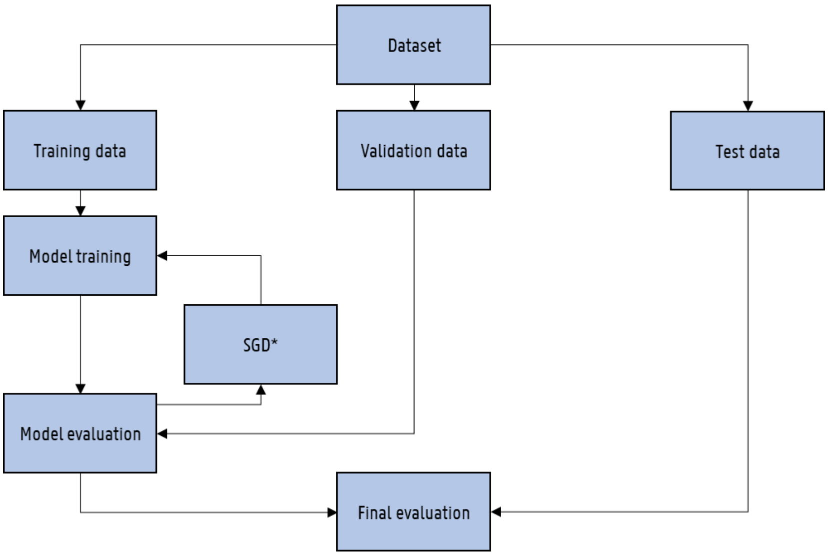

Figure 3: Dataflow of the supervised model

In first instance, the model of choice is YOLOv8. This is a recent and publicly available deep learning model, with a very powerful architecture. The goal of applying this model is to come up with an accurate deep learning model in very little time, so we can already generate weed maps at the entire field level. However, although we will use the same initial architecture for all fields, a specific model will need to be constructed for each LL and crop. The robustness of this model will be verified using consecutive weed datasets, collected later in the season and/or in consecutive years.

Funded by the European Union

Funded by the European Union under Grant Agreement No. 101083589. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or REA. Neither the European Union nor the granting authority can be held responsible for them.

Deep learning is evolving very rapidly, and new and more performant models will become available in the near future. We will therefor follow up and test new models in the project. In addition, once we have a large dataset of all the LL and all present weeds, we will explore the possibility to develop one single model, or to develop a model that is pre-trained on a few datasets, but that can then be trained quickly and with relatively fewer labels using transfer learning.

Unsupervised model workflow

Despite the recent advances in deep learning and AI technology, the supervised models mentioned above still rely on a large number of labels, and are not able to perform weed mapping on datasets of crops/weed combination that the models were not trained for.

Therefore, we will also explore an unsupervised model workflow. The idea behind these models is that the rows, containing the crops, are distinguished automatically from the weeds. With this information, the algorithm can select the plants in-between the rows as the weed plants, and train itself to recognize these plants. Further, we can automatically distinguish different clustering groups (species), so that the model can also recognize these different species automatically. These plants can be classified in order to distinguish the different species. The model is then applied to all plants, also the plants close to the rows, so a full understanding of the weed density per class (species) can be provided.

Model evaluation

The evaluation metrics to be used are based on some of the ones that are used as a standard across most published literature on bounding box object detection. For a given bounding box, the model predicts the location and the confidence.

The confidence can be used as a metric to potentially filter out any weak predictions, but this is only based on the model, and not on any ground-truth. The precise location of the bounding boxes will be used as a metric for the accuracy of the model. Model evaluation on the test dataset, will be used as a general metric to evaluate the performance. On these data, the intersection over union (IoU), precision and recall will be used as general metrics. The formulae for precision and recall can be seen below, a visual representation of IoU is given in Figure 4.

$$Precision = \frac{TP}{TP + FP}$$

$$Recall = \frac{TP}{TP + FN}$$

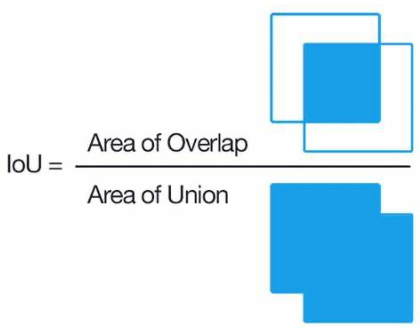

Figure 4: Visual representation of the intersection over union (IoU)

Funded by the European Union

Funded by the European Union under Grant Agreement No. 101083589. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or REA. Neither the European Union nor the granting authority can be held responsible for them.

Insights from AI mapping useful to interpret AWM solutions

Once weed density maps are generated, they will be used throughout the project for several applications:

1) Understanding of treatment levels (cfr supra: cover crop, weed treatment) on weed density per weed species.

   The maps will give an overview of the weed density per species/genus for each individual treatment block. This information is likely to be more precise than the manual weed plot inventories executed. This provides the researchers with a very good opportunity to evaluate the effect of the different treatment on the weed density. Further, once the method has shown to be operational and trustworthy, it can save a lot of time by reducing the inventories.

2) Evolution (dynamics) of weed growth through time per treatment level

   Unique to the GOOD project is that the experiments will be repeated during three consecutive growing seasons. The data are georeferenced, so we can monitor the evolution and dynamics of the different weed species over time.

3) Models as input for site-specific weed management

   The main purpose of the digital maps is that they can eventually be used to automate weed management interventions. This is of course specific for each crop-weed species case, but based on the density maps and the acceptable weed pressure, the field can be divided into blocks, corresponding to the achievable minimal management unit, with each block being assigned the most suitable weed management (e.g., dosage of chemical spraying; necessity to perform weed mechanical intervention or not, …).

Funded by the European Union

Funded by the European Union under Grant Agreement No. 101083589. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or REA. Neither the European Union nor the granting authority can be held responsible for them.