Spatial resolution is often regarded as an important, if not the primary, quality factor of digital images and videos. In many cases, we want to reach beyond the capacities of an imaging sensor, zooming an image more and more. There are a multitude of techniques which are aimed at achieving that goal, and such enhancement process is commonly termed *super-resolution* (SR) reconstruction. It can either be performed from a single image or from many images that show the same scene. The former means that during SR, we are trying to deduce how a given image would look like in higher resolution. This can be done by learning how different objects look like in low and high resolution. For this purpose, the researchers exploit deep learning, in particular the deep *convolutional neural networks* which were found extremely effective in learning complex image analysis pipelines from training data. The deep-learning-based SR solutions allow for generating high-resolution images of natural appearance, even for large zooming factors (e.g. $8\times$). However, as the reconstruction involves some "guessing" on what can be present in the image, the actual high-resolution appearance is not necessarily correctly recovered. An example for that is shown in Figure 1—before reading the figure caption, try to guess which image—(c) or (d) is a real one, and which has been obtained using SR. Although it may be hard to guess which image is real, they would probably be judged as presenting different individuals.

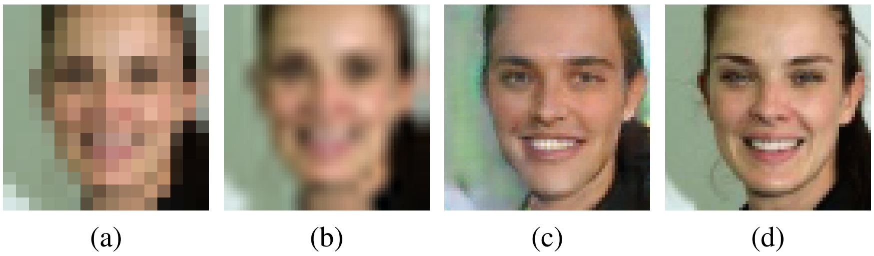

Figure 1: An examples of single-image super-resolution outcome obtained using deep learning for the magnification factor of four: (a) low-resolution image, (b) interpolation outcome, (c) super-resolved image, (d) actual (real) high-resolution image. Source: https://github.com/david-gpu/srez.

Another approach towards SR exploits many images presenting the same scene. Such images may appear similar to each other, but usually there are some tiny differences between them which allow them to contain slightly different portions of information. These differences result from small (sub-pixel) shifts and rotations among the input images, as illustrated in Figure 2. Therefore, these images can be aligned against each other (i.e. co-registered) and combined together to recover the real high-resolution information.

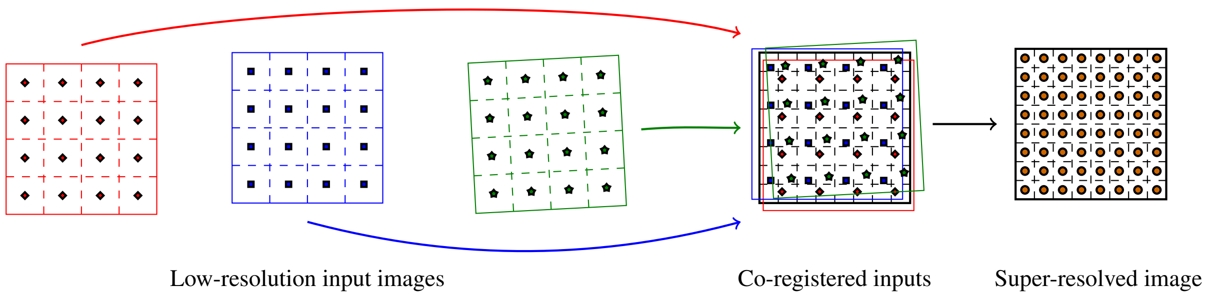

Figure 2: Illustration of information fusion performed from three input images acquired with sub-pixel shifts and rotations. They carry different portions of underlying high-resolution information, which can be exploited during super-resolution reconstruction after co-registration.

The main focus of this project is on developing new algorithms for multiple-image SR aimed at reconstructing the actual high resolution information. We want to combine the benefits of information fusion with the learning capabilities of deep neural networks that are commonly used for single-image SR. In particular, we will introduce new deep architectures for multiple-image SR and propose new ways of representing sets of low-resolution images relying on graph convolutional neural networks. Furthermore, we will develop new algorithms for preparing the training data, and for refining the data presented as an input for reconstruction. Finally, we will propose new loss functions for learning the deep networks that will better reflect the goal of the reconstruction, especially when learning from real-world images. Potential application areas of the developed solutions are concerned with medical imaging, remote sensing, or microscopy imaging.