Computer vision, machine learning, and the promise of phenomics in ecology and evolutionary biology

Moritz Lürig[1]; Seth Donoughe[2]; Erik I. Svensson[1]; Arthur Porto[3,4]; Masahito Tsuboi[1]

[1] Department of Biology, Lund University, 22362 Lund, Sweden
[2] Department of Molecular Genetics and Cell Biology, University of Chicago, Chicago, US
[3] Department of Biological Sciences, Louisiana State University, Baton Rouge, US
[4] Center for Computation & Technology, Louisiana State University, Baton Rouge, US

Abstract

For centuries, ecologists and evolutionary biologists have used images such as drawings, paintings, and photographs to record and quantify the shapes and patterns of life. With the advent of digital imaging, biologists continue to collect image data at an ever-increasing rate. This immense body of data provides insight into a wide range of biological phenomena, including phenotypic trait diversity, population dynamics, mechanisms of divergence and adaptation and evolutionary change. However, the rate of image acquisition frequently outpaces our capacity to manually extract meaningful information from the images. Moreover, manual image analysis is low-throughput, difficult to reproduce, and typically measures only a few traits at a time. This has proven to be an impediment to the growing field of phenomics - the study of many phenotypic dimensions together. Computer vision (CV), the automated extraction and processing of information from digital images, is a way to alleviate this longstanding analytical bottleneck. In this review, we illustrate the capabilities of CV for fast, comprehensive, and reproducible image analysis in ecology and evolution. First, we briefly review phenomics, arguing that ecologists and evolutionary biologists can most effectively capture phenomic-level data by using CV. Next, we describe the primary types of image-based data, and review CV approaches for extracting them (including techniques that entail machine learning and others that do not). We identify common hurdles and pitfalls, and then highlight recent successful implementations of CV in the study of ecology and evolution. Finally, we outline promising future applications for CV in biology. We anticipate that CV will become a

basic component of the biologist's toolkit, further enhancing data quality and quantity, and sparking changes in how empirical ecological and evolutionary research will be conducted.

From phenotypes to phenomics: measuring traits at scale

Faced with the overwhelming complexity of the living world, most life scientists confine their efforts to a small set of observable traits. Although a drastic simplification of organismal complexity, the focus on single phenotypic attributes often provides a tractable, operational approach to understand biological phenomena, e.g. phenotypic trait diversity, population dynamics, mechanisms of divergence and adaptation and evolutionary change. However, there are also obvious limitations in how much we can learn from studying small numbers of *phenotypes* in isolation. Evolutionary and conservation biologist Michael Soulé was one of the first to demonstrate the value of collecting and analyzing many phenotypes at once in his early study of the side-blotched lizard (*Uta stansburiana*; [Soulé 1967]; reviewed in Houle et al.). While doing so, he defined the term "*phenome*" as "the phenotype as a whole" (Soulé 1967). *Phenomics*, by extension, is the comprehensive study of phenomes. In practice, this entails collecting and analyzing multidimensional phenotypes with a wide range of quantitative and high-throughput methods (Houle et al. 2010, Bilder et al. 2009). Given that biologists are now attempting to understand increasingly complex and high dimensional relationships (Walsh 2007), it is surprising that phenomics still remains underutilized (Fig. 1), both as methodological approach and as an overarching conceptual and analytical framework (Houle et al. 2010).

Phenomic datasets are essential if we are to understand some of the most compelling but challenging questions in the study of ecology and evolution. For instance, natural selection typically does not operate on single traits, but on multiple traits simultaneously (Lande and Arnold 1983, Phillips and Arnold 1999b). Such correlational selection can bias evolutionary change and shape genetic covariance patterns by building up and maintaining linkage disequilibrium (Schluter 1996, Phillips and Arnold 1999a, Sinervo and Svensson 2002,

Svensson et al. 2021), and it is a phenomenon that can most effectively be uncovered in multidimensional phenotypic datasets. Another example is pleiotropy, which generates patterns of covariation among traits that are impossible to predict if only a few simple traits are measured (Visscher and Yang 2016, Saltz et al. 2017). Phenotypic plasticity, which is increasingly recognized in mediating evolutionary trajectories (Pfennig et al. 2010), is also an inherently multivariate phenomenon involving many traits and interactions between traits, so it should be quantified as such (Morel-Journel et al. 2020). Finally, community stability depends on species interactions and ecological niches of organisms; niches are typically influenced by many traits at once (Blonder 2018, Laughlin et al. 2020). Put simply: if we are to draw a complete picture of biological processes and aim to understand their causal relationships at various levels of biological organization, we need to measure more traits. Phenomic datasets will make our conclusions and inferences more robust if underpinned by more complete information without systematic biases.

High dimensional phenotypic data are also needed for uncovering the causal links between genotypes, environmental factors, and phenotypes, i.e. to understand the genotype-phenotype map (Houle et al. 2010, Orgogozo et al. 2015). The advent of genomics - high throughput molecular methods to analyze the structure, function or evolution of an organism's genome in parts or as a whole (Church and Gilbert 1984, Feder and Mitchell-Olds 2003) - has already improved our understanding of many biological phenomena. This includes the emergence and maintenance of biological diversity (Seehausen et al. 2014), the inheritance and evolution of complex traits (Pitchers et al. 2019), and the evolutionary origin of key metabolic traits (Ishikawa et al. 2019). Thus, accessible molecular tools have lowered the hurdles for discovery-based genomic research and shifted the focus away from the study of observable organismal traits and phenotypes towards their molecular basis. However, a similar “moonshot-program” for the phenotype, i.e. an ensemble of phenomics methods that matches genomics in their comprehensiveness, is still lacking (Freimer and Sabatti 2003). The growing mismatch in how efficiently molecular and phenotypic data are collected may hamper

82 further scientific progress in ecological and evolutionary research (Houle et al. 2010, Orgogozo
83 et al. 2015, Lamichhaney et al. 2019).

84 Following previous calls for phenomic research programs (Bilder et al. 2009, Houle et
85 al. 2010, Furbank and Tester 2011), some recent studies have collected phenotypic data with
86 high dimensionality, for example, in plants (Ubbens and Stavness 2017), animals (Cheng et
87 al. 2011, Kühl and Burghardt 2013, Pitchers et al. 2019) and microbes (Zackrisson et al. 2016,
88 French et al. 2018). The methods in such studies included 2D- and 3D-scanners, camera
89 traps, robotic imaging platforms, computational image analysis, morphometrics,
90 transcriptomics, metabolomics, and automated data loggers to record physiological and
91 behavioral data from organisms (Houle et al. 2010). Many of these techniques produce image-
92 based data. In general, ecologists and evolutionary biologists use digital imaging to quantify
93 the external phenotype of an organism (i.e. its visually observable characteristics), to count
94 organisms (e.g. cells on microscope slides), or to detect the presence of an organism (e.g. in
95 images collected by camera traps). Existing work has supplied us with an immense body of
96 image data that has provided insight into a wide range of biological phenomena, yet, when
97 biologists manually extract phenotypes from images for phenomic-scale research, they
98 confront several main bottlenecks.

99 A major constraint when working with large amounts of images (~1000 or more) is
100 processing time and cost. Manual extraction of phenotypic data from images is slow and it
101 requires trained domain experts whose work is extremely expensive. Moreover, the collection
102 of such metrics in a manual fashion entails subjective decisions by the researcher, which may
103 make it prone to error, and certainly makes reproducibility difficult. Last, manually measured
104 traits tend to be low-dimensional measurements of higher dimensional traits. For example,
105 external color traits, such as eye color phenotypes, are often scored as discrete categories
106 (e.g. red vs blue phenotypes), whereas *pixel* level information (number of red vs. blue pixels)
107 can provide a continuous phenotypic metric (Liu et al. 2010). Such quantitative, high-
108 dimensional data can provide insight into previously hidden axes of variation in natural
109 phenotypes. In this review we extol *computer vision* (CV), the automatic extraction of

meaningful information from images, as a promising toolbox to collect phenotypic information on a massive scale. The field has blossomed in recent years, producing a diverse array of computational tools to increase analytic efficiency, data dimensionality, and reproducibility. We argue that CV is poised to become a basic component of the data analysis toolkit in ecology and evolution, enabling researchers to collect and explore phenomic-scale data.

Digital images as data

The structure of digital images

A two-dimensional image is an intuitive way to record, store, and analyze organismal phenotypes. In the pre-photography era, ecologists and evolutionary biologists used drawings to capture the shapes and patterns of life, later to be replaced by analog photography, which allowed for qualitative assessment and simple, often only qualitative analysis of phenotypic variation. With the advent of digital photography, biologists could collect phenotypic data at unprecedented rates using camera stands, camera traps, microscopes, scanners, video cameras, or any other instrument with semiconductor image sensors (hereafter “image sensors”). Image sensors produce two-dimensional raster images (also known as bitmap images), which store incoming visible light or other electromagnetic signals into discrete, locatable picture elements - in short: pixels (Fig. 2). Each pixel contains quantitative phenotypic information that is organized as an array of rows and columns, whose dimensions are also referred to as “pixel resolution” or just “resolution”. An image with 1000 rows and 1500 columns has a resolution of 1000 x 1500 (= 1 500 000 pixels, or 1.5 megapixels). The same applies for digital videos, which are simply a series of digital images displayed in succession, where the frame rate (measured as frames per second = fps) describes the speed of that succession.

At the pixel level, images or video frames can store variable amounts of information, depending on the *bit depth*, which refers to the number of distinct values that a pixel can represent (Fig. 2). In binary images, pixels contain information as a single bit, which can take

exactly two values - typically black or white (2^1 values = 2 intensity values). Grayscale images from typical consumer cameras have a bit depth of 8, thus each pixel can take a value between 0-255 (2^8 values = 256 intensity values), which typically represents a level of light intensity, also referred to as pixel intensity. Color images are typically composed of at least three sets of pixel arrays, also referred to as channels, each of which contain values for either red, green or blue (RGB; Fig. 2). Each channel, when extracted from an RGB image, is a grayscale representation of the intensities for a single-color channel. Through the combination of pixel values at each location into triplets, colors are numerically represented. Today the industrial standard for color images is 24-bit depth, in which each color channel has a bit depth of 8 and thus can represent 256 colors (Fig. 1). Thus, 24-bit RGB images can represent over 16 million color variations in each pixel (2^{24} = 256 ^{x} 256 ^{x} 256 = 16 777 216 intensity values), which already greatly surpasses the estimated 2.28 million of color variations that humans can perceive (Pointer and Attridge 1998).

Today, high resolution image sensors are an affordable way to store externally visible phenotypic information, like color and shape. However, advanced image sensors can also combine information from different spectra other than the visible light, like infrared radiation, which can be used to quantify individual body temperatures. With thermal image sensors, biologists can estimate body surface temperatures, which are correlated with internal (core) body temperatures (Tattersall and Cadena 2010), particularly in small animals like insects (Tsubaki et al. 2010, Svensson et al. 2020). Thermal imaging, or thermography, offers new opportunities for ecophysiological evolutionary research of how animals cope with heat or cold stress in their natural environments (Fig. 4B; Tattersall et al. 2009, Tattersall and Cadena 2010, Svensson and Waller 2013). Fluorescence spectroscopy is another way to quantify phenotypes in high throughput and with high detail. For example, plate readers typically used in microbial and plankton research, can combine light in the visible spectrum with images containing information of cell fluorescence or absorbance to an “image stack” (Roeder et al. 2012). Image stacks and the inclusion of multiple spectral channels provide a promising

avenue of research towards capturing a more complete representation of the phenotype (Fig. 4A; Hense et al. 2008, Di et al. 2014).

A brief introduction to computer vision

CV is an interdisciplinary field at the intersection of *signal processing* and *machine learning* (Fig. 6; Mitchell 1997), which is concerned with the automatic and semiautomatic extraction of information from digital images. After image acquisition, a typical CV workflow involves three main steps: preprocessing, *segmentation*, and measurement (Fig. 3).

Preprocessing - Independent of how much care has been taken during image acquisition, preprocessing is an important step to prepare images for the CV routines to follow. There is a wealth of image processing techniques that can be applied at this stage, such as transformations to reduce noise (e.g. gaussian blur) or enhance contrast (e.g. histogram adjustment). Images can also be masked or labeled to filter the image so that subsequent steps are applied to the intended portions of each image. Defining the appropriate coordinate space (i.e. pixel-to-mm ratios) is also part of preprocessing. This step is highly specific to the respective study system or image dataset and may initially require some fine-tuning by the scientist to ensure data quality, which, however, can typically be automated afterwards.

Segmentation The central step in all phenotyping related CV-pipelines is the segmentation of images into pixels that contain the desired signal or information (*foreground*) and all other pixels (background). In its most basic form, segmentation of grayscale images can be done by simple signal processing algorithms, such as a *threshold* or *watershed*. Similarly, *feature detection* algorithms examine pixels and their adjacent region for specific characteristics or key points, e.g. whether groups of pixel form edges, corners, ridges or blobs. Videos or multiple images of the same scene provide an additional opportunity for segmentation: foreground detection can detect changes in image sequences to determine the pixels of interest (e.g. a specimen placed in an arena, or animals moving against a static background), while subsequent background subtraction isolates the foreground for further processing. Finally, *object detection* describes the high-level task of finding instances of *semantic objects*

(organisms, organs, structures, etc.) in an image, which is typically addressed through classical machine learning or *deep learning* (see section Computer vision methods). In classical machine learning, *features* must be first engineered or extracted from a training dataset using *feature detectors*, then used to train a classifier, and finally applied to the actual dataset (Mitchell 1997). Deep learning algorithms are a family of machine learning methods based on artificial *neural networks* that "learn" what constitutes the object of interest during the training phase. With sufficient training from *labelled* images, deep learning-powered object detection algorithms can be highly accurate and often greatly outperform pre-existing object recognition methods - in some cases even human experts (Buetti-Dinh et al. 2019).

Measurement What sorts of data can we extract from images? CV can retrieve a multitude of phenotypic traits from digital images in a systematic and repeatable fashion. In the simplest case such traits have been measured traditionally and are established in each study system, such as body size (e.g. length or diameter) or color (e.g. "red phenotype" vs "blue phenotype"). In such cases, switching from a manual approach to a semi- or fully automatic CV approach is straightforward, because the target traits are well embedded in existing statistical and conceptual frameworks. The main benefits from CV are that costly manual labor is reduced and that the obtained data becomes more reproducible, because the applied CV analysis pipeline can be stored and re-executed. It is also possible to increase the number of dimensions without much extra effort and without discarding the traditionally measured traits (Table 1). For example, in addition to body size, one could extract body shape traits, i.e. the outline of the body itself (i.e. contour coordinates of the foreground), and texture (i.e. all pixel intensities within the foreground). Such high dimensional traits can be directly analyzed using multivariate statistics, or transformed into continuous low dimensional traits, such as continuous shape features (circularity or area), texture features (color intensity or variation, pixel distribution), or moments of the raw data (Table 1).

A history of computer vision methods

The field of CV is now close to celebrating its 6th decade. It first emerged in the late 1950s and early 1960s, in the context of artificial intelligence research (Rosenblatt 1958). At the time, it was widely considered a stepping-stone in our search for understanding human intelligence (Minsky 1961). Given its long history, a wide variety of CV techniques have emerged since its inception, but they all contain variations of the same basic mechanism. CV is, from the methodological standpoint, the process of extracting meaningful features from image data and then the use of such features to perform tasks, which, as described above, may include classification, segmentation, recognition, detection, among others. In this section, we will not aim at presenting an all-encompassing review of all CV methods, but rather to identify the major trends in the field and highlight the techniques that have proved useful in the context of biological research. It is worth noting that even classical CV approaches are still routinely used in the modern literature, either in isolation or, most commonly, in combination with others. In a large part, methodological choices in CV are highly domain-specific (see section *Practical considerations for computer vision,* Fig. 4, and Fig. 6).

First wave - Hand-crafted features

The first wave of CV algorithms is also the closest one to the essence of CV, namely, the process of extracting features from images. Starting with the work of Larry Roberts, which aimed at deriving 3D information from 2D images (Roberts 1963), researchers in the 1970s and 1980s developed different ways to perform feature extraction from raw pixel data. Such features tended to be low-level features, such as lines, edges, texture or lighting, but provided us with the initial basic geometric understanding of the data contained in images. A notable example of such algorithms is the watershed algorithm. First developed in 1979 (Beucher 1979), the watershed algorithm treats images as a topographic map, in which pixel intensity represents its height, and attempts to segment the image into multiple separate ‘drainage basins’. This algorithm is still routinely used in signal processing techniques (Fig. 6) and can

be effectively used to process biological images such as those obtained through animal or plant cell microscopy (McQuin et al. 2018). Other initial low-level hand-crafted approaches that achieved popularity include the Canny and Sobel filters (edge detectors; Canny 1986, Kanopoulos et al. 1988) and Hough transforms (ridge detection; Duda and Hart 1972).

Another approach that gained popularity in the CV literature in the early 1990s was principal component analysis (PCA). In a PCA, independent, aggregate statistical features are extracted from multidimensional datasets. These can be used, for example, in classification. One of the most notable uses of PCA in the context of CV was the eigenfaces approach (Turk and Pentland 1991). Essentially, Turk and Pentland (1991) noted that one could decompose a database of face images into eigenvectors (or characteristic images) through PCA. These eigenvectors could then be linearly combined to reconstruct any image in the original dataset. A new face could be decomposed into statistical features and further compared to other known images in a multidimensional space. PCA has notably found many other uses in biology (e.g., (Ringnér 2008).

In the late 1990s and early 2000s, Scale Invariant Feature Transform (SIFT; Lowe 1999, 2004) and Histogram of Oriented Gradients (HOG; Dalal and Triggs 2005) were developed. Both SIFT and HOG represent intermediate-level local features that can be used to identify keypoints that are shared across images. In both approaches, the first step is the extraction of these intermediate-level features from image data, followed by a feature matching step that tries to identify those features in multiple images. Finding keypoints across images is an essential step in many CV applications, such as object detection, landmarking, and image registration. These intermediate-level features have several advantages over the lower-level features mentioned above, most notably the ability to be detected in a wide-variety of scales, noise and illumination. Another key aspect of SIFT and HOG features is that they are generally invariant to certain geometric transformations, such as uniform scaling and simple affine distortions.

Second wave - Initial machine-learning approaches

While the use of hand-crafted features spurred much of the initial work in CV, soon it became apparent that without image standardization, those low- and intermediate-level features will often fall short of producing sufficiently robust CV algorithms. For example, images belonging to the same class can often look very different and the identification of a common set of shared low-level features can prove to be quite challenging. Consider, for instance, the task of classifying animal images. Two cat breeds can look quite different, despite belonging to the same general class (cat). As such, while the initial feature-engineering approaches were essential for the development of the field, it was only with the advent of machine-learning that CV acquired more generalizable applications.

Machine learning algorithms for CV can be divided in two main categories (but see Box 2): supervised and unsupervised (Geoffrey Hinton, Terrence J. Sejnowski 1999). Unsupervised algorithms attempt to identify previously unidentified patterns on unlabeled data. In other words, no supervision is applied to the algorithm during learning. While it can be argued that PCA was one of the first successful unsupervised learning algorithms applied directly to CV, here we group PCA with “first wave” tools due to its use as a feature extractor. Other unsupervised learning algorithms commonly used in CV include clustering techniques, such as k-means (Lloyd 1982) and gaussian mixture models (GMM; Reynolds and Rose 1995). Clustering algorithms represented some of the first machine learning approaches for CV. Their aim is to find an optimal set of objects (or components) that are more similar to each other than to those in other sets. This type of approach allowed researchers to find hidden patterns embedded in multidimensional data, proving useful for classification and segmentation tasks.

However, it is in the supervised domain that machine learning for CV has been most successful (Heileman and Myler 1989). In supervised learning approaches, the user supplies labeled training data in the form of input-output pairs (Box 2). The ML algorithm iteratively “learns” a function that maps input into output for the labeled training data. Among the initial supervised learning approaches for CV, Support Vector Machines (SVM) were by far the most

common approach (Cortes and Vapnik 1995). Given a certain image dataset and their corresponding labels (e.g., classes in a classification task), SVMs find the hyperplane (in feature space) that maximizes the separation between the classes of interest. An essential aspect of SVMs is that such learned decision boundaries separating the classes can be nonlinear in the original feature space, allowing the model to separate classes that would not be separable by a purely linear technique (Cortes and Vapnik 1995).

Third wave - Ensemble methods

While SVMs were extremely successful in CV and spurred much of the supervised work that happened afterwards, it became clear by the early 2000s that single estimators often underperformed approaches combining the predictions of several independent estimators, an approach known as ensemble methods (Dietterich 2000). Ensemble methods represent a slightly different philosophical approach to machine learning, in which multiple models are trained to solve the same task and their individual results are combined to obtain an even better model performance. Several ensemble methods have been developed in the literature, but they are generally divided in two main families: bagging and boosting.

Bagging approaches combine several models that were trained in parallel through an averaging process (Bauer and Kohavi 1999). Each underlying model is trained independently of the others based on a bootstrap resample of the original dataset. As a consequence, each model is trained with slightly different and (almost) independent data, greatly reducing the variance in the combined model predictions. A classic example of bagging approach is the random forest algorithm (Breiman 2001), in which multiple learning trees are fitted to bootstrap resamples of the data and posteriorly combined through mean averaging (or majority vote).

Boosting, on the other hand, combines learners sequentially rather than in parallel (Bauer and Kohavi 1999). Among boosting algorithms, gradient boosting (Friedman 2000) is one of the most widely used in CV. In gradient boosting, models are combined in a cascade fashion, such that a downstream model is fitted to the residuals of upstream models. As a consequence, while each individual model in the cascade is only weakly related to the overall

task, the combined algorithm (i.e., the entire cascade) represents a strong learner that is directly related to the task of interest. Since this approach, if unchecked, will lead the final model to overfit the training data, regularization procedures are usually applied when using gradient boosting.

Fourth wave – Deep learning

Deep learning approaches are, at the time of this writing, the state-of-the-art in CV and have recently become more accessible through the community-wide adoption of code-sharing practices (e.g. via https://github.com/ or https://stackoverflow.com/). *Deep learning* refers to a family of machine learning methods based on artificial neural networks with multiple steps that perform *convolutions* or other mathematical operations on the input data, each of which is referred to as a *hidden layer.* Networks with dozens or hundreds of hidden layers (i.e. deep networks) allow for the extraction of high-level features from raw image data (LeCun et al. 2015). While they have only recently become widespread, the history of artificial neural networks is at least as old as the field of CV itself. One of first successful attempts in the study of artificial neural networks was the perceptron (Rosenblatt 1958), a computer whose hardware design was inspired by neurons, and which was used to classify a set of inputs into two categories. This early work, while successful, was largely restricted to linear functions and therefore could not deal with non-linearity, such as XOR functions (Minsky et al. 1969). As a consequence, artificial neural network research remained rather understudied until the early 80s when training procedures for multi-layer perceptrons were introduced (i.e., backpropagation; Rumelhart and McClelland 1987). Even then, multi-layer approaches were computationally taxing, and the hardware requirements represented an important bottleneck to research in neural network-based CV, which remained disfavored compared to much lighter approaches, such as SVMs.

It eventually became clear that a neural network approach to CV represented a fundamental leap for CV. When compared to the hand-crafted features that dominated the field for most of its history, neural networks learn features from the data itself, therefore

eliminating the need for feature engineering (LeCun et al. 2015). In a large part, deep learning approaches for CV have only emerged in force due to two major developments at the beginning of the 21st century. On one side, hardware capability greatly increased due to high consumer demand for personal computing and gaming. On the other, there was a widespread adoption of the internet, leading to an exponential increase in data availability through shared image databases and labelled data. Today, deep learning is a general term that encompasses a wide variety of approaches that share an architectural commonality of relying on training neural networks with multiple hidden layers. However, this superficial similarity hides a considerable array of differences between different algorithms and one could say that the field of deep learning is as diverse as the domains in which CV is applied. We present some of the most relevant classes of deep learning approaches in Box 2.

Practical considerations for computer vision

Measurement theory: define your traits thoughtfully

Defining meaningful phenotypes is deceptively challenging. Traditionally, biologists relied on intuition and natural history conventions to define phenotypes, but this approach can obscure the fact that phenomes are exceedingly high-dimensional, and many dimensions are infinitely divisible. When deciding what to measure, we suggest that researchers consider *measurement theory,* a qualitative formalization of the relationship between actual measurements and the entity that the measurements are intended to represent (Houle et al. 2011). In phenomics using CV, we recommend that researchers adhere to the following three principles: i) Ensure that the measurements are meaningful in the theoretical context of research questions. ii) Remember that all measurements are estimates. Measurements without uncertainties should always be avoided. iii) Be careful with units and scale types, particularly when composite values, such as the proportion of one measurement over another, are used as a measurement. Wolman (2006) and Houle et al. (2011) give details of measurement theory and practical guidelines for its use in ecology and evolutionary biology.

Image quality and pertinent metadata: collect images that are maximally useful

As a general rule of thumb, images taken for any CV analysis should have a *signal-to-noise* ratio (SNR) sufficiently high so that the signal (i.e. the phenotypic information) is detectable from the image background. High SNR can be achieved by using high resolution imaging devices (e.g. DSLR cameras or flatbed scanners), ensuring that the object is in focus (e.g. automatically or by fixing the distance between camera and object), and by creating a high contrast between object and background (e.g. by using backgrounds that are of contrasting color or brightness to the organism or object). We recommend to iteratively assess suitability of imaging data early on in a project and adjust if necessary. This means taking pilot datasets, processing them, measuring traits, estimating measurement errors, and then updating the image collection process. Moreover, it is good practice to include a color or size reference whenever possible. It helps researchers to assess if the image has sufficient SNR, increases reproducibility, and helps to evaluate measurement bias as we discuss in the next section.

Every measurement is an estimate: dealing with measurement error

Another important aspect to consider when measuring phenotypes from images is the (in)accuracy of measurements. Formally, measurement inaccuracy is composed of imprecision and bias, corresponding to random and systematic differences between measured and true values, and can be expressed as the following relationship

$$\text{inaccuracy} = \text{imprecision} + \text{bias}^2$$

(Grabowski and Porto 2017, Tsuboi et al. 2020). These two sources of errors characterize distinct aspects of a measurement: precise measurements may still be inaccurate if biased, and unbiased measurements may still be inaccurate if imprecise (Fig. 5). Measurement imprecision can be evaluated by the coefficient of variation (standard deviation divided by the mean) of repeated measurements. Bias requires a knowledge of true values.

We ultimately need to understand if a measurement is sufficiently accurate to address the research question at hand. Repeatability is a widely used estimator of measurement accuracy in ecology and evolutionary biology (Wolak et al. 2012), which in our notation could be expressed as

$$\text{repeatability} = 1 - \frac{inaccuracy}{total\ variance}$$

This expression clarifies that the repeatability depends both on measurement inaccuracy and total variance in the data. For example, volume estimates of deer antler from 3D photogrammetry have an average inaccuracy of 8.5%, which results in repeatabilities of 67.8-99.7% depending on the variance in antler volume that a dataset contains (Tsuboi et al. 2020). In other words, a dataset with little variation requires more accurate measurement to achieve the same repeatability as a dataset with more variation. Therefore, the impact of measurement error has to be evaluated in the specific context of data analysis.

CV-based phenomics is extremely useful in this regard because it allows researchers to identically repeat a measurement process, and thereby evaluate inaccuracies in order to improve measurement precision. In the aforementioned example of volume estimated from 3D photogrammetry (Tsuboi et al. 2020), it was found that 70% of the total inaccuracy arose from the error in scaling arbitrary voxel units into real volumetric units. Therefore, by using the mean of two estimates obtained from two copies of an image that are scaled twice independently as a representative measurement, the inaccuracy dropped to 5.5%. However, the opportunity to improve accuracy by repeated measurements is limited if a majority of error arises from the stored images themselves. For this reason, we recommend always taking repeated images of the same subject at least for a subset of data. This will allow evaluating the magnitude of error due to images relative to the error due to acquisition of measurements from images. If the error caused by images is large compared to the error caused by data

acquisition, it may be necessary to modify imaging and/or preprocessing protocol to increase SNR.

Assessing measurement bias requires separate treatments. When linear (length) or chromatic (color) measurements are obtained from images, it is a good general practice to include size and color scales as part of images to estimate bias as the difference between known values of imaged scales and measurements obtained through CV (i.e. the reference card in Fig. 3). Knowing the true value may be difficult in some cases, such as domain area or circularity (Hoffmann et al. 2018), since they are hard to characterize without a CV. When multiple independent methods to measure the same character exist, we recommend using them on sample data to determine the bias of one method relative to the other.

Selecting a CV pipeline: as simple as possible, as complex as necessary

When using CV tools there are usually many different ways to collect a specific type of phenotypic information from images (Fig. 6). Therefore, one of the first hurdles to overcome when considering the use of CV is selecting the appropriate technique from among a large and growing set of choices. The continued emergence of novel algorithms to collect, process and analyze image-derived data may sometimes make us believe that any “older” technology is immediately outdated. Deep learning, specifically CNNs, is a prominent example of an innovation in CV that was frequently communicated as so “revolutionary” and “transformative” that many scientists believed it would replace all existing methods. However, despite the success of CNNs, there are many cases where they are inappropriate or unfeasible, e.g. due to small sample sizes, hardware or time constraints, or because of the complexity that deep learning implementations entail, despite many efforts to make this technology more tractable. We discourage readers from defaulting to using the newest technology stacks; rather, we suggest that researchers be pragmatic as to which is the fastest and simplest way to get the phenotypic information of desire from any given set of images.

Begin by considering the size of a given image dataset, whether it is complete, e.g. after an experiment, or whether there will be continued future additions, e.g. as part of a long-

462 term experiment or field survey. As a rough rule of thumb, if a dataset encompasses only a
463 thousand images or fewer, consider it “small”; if a dataset has thousands to tens of thousands
464 images, consider it “large” (see Fig. 6 for methodological suggestions for each case). The next
465 assessment should be about the SNR in your images: images taken in the laboratory typically
466 have a high degree of standardization, e.g. controlled light environment or background, and
467 thus a very high SNR. Field images can also have a high SNR, for example, if they are taken
468 against the sky or if the trait of question is very distinct from the background through bright
469 coloration. If the dataset is “small” and/or has high SNR, it may not be necessary to use the
470 more sophisticated CV tools; instead, signal processing, e.g. threshold or watershed
471 algorithms, may already be sufficient for segmentation although typically some pre- and post-
472 processing is typically still required (e.g. blurring to remove noise, “morphology”-operations to
473 close gaps, or masking false positives).
474 For large datasets, images with low SNR, or if the information of interest is variable
475 across images (e.g. traits are photographed from different angles or partially covered up),
476 machine learning approaches are probably more suitable. In contrast to signal processing,
477 where segmentation results are immediately available, all machine learning image analysis
478 pipelines include iterative training and validation phases, followed by a final testing phase.
479 Such a workflow can be complex to initiate but pays off in the long run by providing
480 segmentation results that become increasingly robust if more training data is supplied over
481 time. Classic machine learning algorithms often require an intermediate amount of training
482 data (500-1000 or more images) before they can produce satisfactory results. In this category,
483 SVM or HOG algorithms are a good choice when areas of interest do not contrast sufficiently
484 from the surrounding area, for example, when automatically detecting landmarks (Porto and
485 Lysne Voje 2020). Deep learning algorithms require much larger training datasets (several
486 1000s to 10000s) but are less sensitive to noise and idiosyncrasies of the foreground. Thus,
487 for large and continuously growing data sets, or for recurring image analysis tasks, deep
488 learning has become the standard approach for segmentation. Deeper networks may
489 increase model accuracy, and thus improve the segmentation results, but have an increasing

risk of overfitting the contained information - i.e. the model is less generalizable to input data. While the implementation of deep learning pipelines requires some expertise, they can be retrained and are typically less domain specific than classic machine learning pipelines (O'Mahony et al. 2020).

Recent examples of computer vision to collect phenomic data

"Phenomics" as a term has not yet gained widespread attention in the ecological and evolutionary biology research communities (Fig. 1), but many biologists are engaged in research programs that are collecting phenomic data, even though it is not called as such. Some of them are already using automatic or semi-automatic CV to collect phenotypic data. Here we present small a selection of promising applications of CV to answer ecological or evolutionary research questions (points matching panels in Fig. 4):

A. **Resource competition traits -** Species diversity within ecological communities is often thought to be governed by competition for limiting resources (Chesson 2000). However, the exact traits that make species or individuals the best competitors under resource limitation conditions are difficult to identify among all other traits. In this example, the phenotypic space underlying resource competition was explored by implementing different limitation scenarios for experimental phytoplankton communities. Images were taken with a plate reader that used a combination of visible light and fluorometry measurements (Hense et al. 2008). The images were analyzed using signal processing, which allowed the rapid segmentation of several 1000 images by combining information from multiple fluorescence emission excitation spectra to an image stack. As a result, over 100 traits related to morphology (shape, size, and texture) and internal physiology (pigment content, distribution of pigments within each cell) were obtained at the individual cell level. (Gallego et al., unpublished data)

B. **Thermal adaptation and thermal reaction norms** - Variation in body temperature can be an important source of fitness variation (Kingsolver and Huey 2008, Svensson

et al. 2020). Quantifying body temperature and thermal reaction norms in response to natural and sexual selection allows us to test predictions from evolutionary theory about phenotypic plasticity and canalization (Lande 2009, Chevin et al. 2010). However, body temperature is an internal physiological trait that is difficult to quantify in a non-invasive way on many individuals simultaneously and under natural conditions. Thermal imaging is an efficient and non-invasive method to quantify such physiological phenotypes on a large scale and can be combined with thermal loggers to measure local thermal environmental conditions in the field (Svensson and Waller 2013, Svensson et al. 2020).

C. **Stochastically patterned morphological traits** - In contrast to homologous, landmark-based morphological traits, tissues also form emergent patterns that are unique to every individual. The arrangement of veins on the wings of damselflies is one such example. By measuring the spacing, angles, and connectivities within the adult wing tissue, researchers have proposed hypotheses about the mechanisms of wing development and physical constraints on wing evolution (Hoffmann et al. 2018, Salcedo et al. 2019).

D. **Morphometrics and shape of complex structures** - Landmark-based morphometrics has become a popular tool used to characterize morphological variation in complex biological structures. Despite its popularity, landmark data is still collected mainly through manual annotation, a process which represents a significant bottleneck for phenomic studies. However, machine-learning-based CV can be used to accurately automate landmark data collection in morphometric studies not only in 2D (McPeek et al. 2008, Porto and Voje 2020), but also in 3D (Porto et al. 2020).

E. **Volumes of morphologically complex traits**. Many topics in evolutionary ecology concerns investment of resources into a particular trait. However, measuring energetic investment, either as mass or volume of the target traits, has been challenging because many traits are morphologically complex, making it difficult to estimate investment from a combination of linear measurements. Photogrammetry is a low-cost

and fast technique to create 3D surface images from a set of images. Using a simple protocol and a free proprietary software, Tsuboi et al. (2020) demonstrated that photogrammetry can accurately measure the volume of antler in deer family Cervidae. The protocol is still relatively low-throughput due primarily to the necessity of high number of images (> 50) per sample, but it allows extensive sampling (*sensu* Houle et al. 2010) of linear, area and volumetric measurements of antler structures.

Outlook

In this review we provided a broad overview of various CV techniques and gave some recent examples of their application in ecological and evolutionary research. We presented CV as a promising toolkit to overcome the image analysis bottleneck in phenomics. However, to be clear, we do not suggest that biologists discontinue the collection of univariate traits like body size or discrete colors. Such measures are undoubtedly useful if they contain explanatory value and predictive power. Instead, we propose that CV can help to i) collect them with higher throughput, ii) in a more reproducible fashion, and to iii) collect additional traits so we can interpret them in the context of trait combinations. We argue that CV is not bound to immediately replace existing methods, but it simply opens the opportunity to place empirical research of phenotypes on a broader base. We also note that CV based phenomics can be pursued in a deductive or inductive fashion. In the former case, scientists would simply conduct hypothesis driven research including a wider array of traits into causal models (Houle et al. 2011); in the latter, they would engage in discovery-based data-mining approaches that allow scientists to form hypotheses a posteriori based on the collected data (Kell and Oliver 2004).

Although CV based phenomics provides new opportunities for many areas of study, we identify several specific fields that will profit most immediately from CV. First, evolutionary quantitative genetics will benefit tremendously from increased sample sizes that CV-phenomics entails, because the bottleneck of the field has been the difficulty in accurately

570 estimating key statistics such as genetic variance covariance matrices and selection gradients.
571 The recent discovery of tight matches between mutational, genetic, and macroevolutionary
572 variances in Drosophilid wing shape (Houle et al. 2017) is exemplary of a successful phenomic
573 project. Second, large-scale empirical studies of the genotype-phenotype map will finally
574 become possible, because of the availability of high-throughput phenotypic data and analytical
575 framework to deal with big data (Pitchers et al. 2019, Zheng et al. 2019, Maeda et al. 2020).
576 Third, studies of fossil time-series will gain opportunities to document and analyze the
577 dynamics of long-term phenotypic evolution with unprecedented temporal resolution (Liow et
578 al. 2017, Brombacher et al. 2017). The ever-growing technology of CV indicates that these
579 are likely a small subset of unforeseen future applications of CV phenomics in our field. Just
580 like the technological advancements in DNA-sequencing has revolutionized our view of
581 genomes, development and molecular evolution in the past decades, we anticipate that the
582 way we look at phenotypic data will be changing in the coming years.

583 Just as CV is changing what it means to measure a trait, there is a complementary
584 change in what can be considered scientific image data in the first place. Large, publicly
585 available image datasets are fertile ground for ecology and evolutionary research. Such
586 databases include both popular and non-scientific social media (e.g. https://www.flickr.com/
587 or https://www.instagram.com/), but also quality-controlled and vetted natural history and
588 species identification resources with global scope and ambitions (e.g.
589 https://www.inaturalist.org/). Successful examples of how such public image databases can
590 be useful are in studies aiming to quantify the frequencies variation of discrete traits, such as
591 color polymorphism frequencies in different geographic regions (Leighton et al. 2016). These
592 manual efforts in mining available public image resources can potentially be replaced in the
593 future using more automated machine learning or CV approaches. Similarly, the corpus of
594 published scientific literature is full of image data that can be combined and re-analyzed in
595 order to address larger-scale questions (Hoffmann et al. 2018, Church et al. 2019a, 2019b).

596 Previous calls for phenomics argued that, to make phenomics a successful endeavor,
597 it has to be extensive, aiming at measuring many different aspects of the phenotypes, as well

as intensive, aiming at characterizing each measurement accurately with large sample size and with high temporal resolution (Bilder et al. 2009, Houle et al. 2010, Furbank and Tester 2011). We agree with this view, but we also emphasize that phenomics is nothing conceptually new in this respect. As we discussed, many researchers in our field have already adopted phenomic pipelines, studying high-dimensional phenotypic data acquired by high-throughput measuring technologies without using the term phenomics. If so, what is the conceptual value of phenomics? In our opinion, phenomics is more than just a rigorous version of conventional research of organismal phenotypes, but also a dedication towards phenotypic data. Phenomics shifts us from viewing phenotypes as given entities towards viewing them as part of the phenome at the whole organismal level.

Acknowledgements

The publication of this study was funded through the Swedish Research Council International Postdoc Grant (2016-06635) to MT. ML was supported by a Swiss National Science Foundation Early Postdoc.Mobility grant (SNSF: P2EZP3_191804). EIS was funded by a grant from the Swedish Research Council (VR: grant no. 2016-03356). SD was supported by the Jane Coffin Childs Memorial Fund.

Authors contributions

ML conceived the idea for this review and initiated its writing. In the process, all authors contributed equally to the development and discussion of ideas, and to the writing of the manuscript.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

Bauer, E., and R. Kohavi. 1999. An Empirical Comparison of Voting Classification Algorithms: Bagging, Boosting, and Variants. Machine learning 36:105–139.

Beucher, S. 1979. Use of watersheds in contour detection. Proceedings of the International Workshop on Image.

Bilder, R. M., F. W. Sabb, T. D. Cannon, E. D. London, J. D. Jentsch, D. S. Parker, R. A. Poldrack, C. Evans, and N. B. Freimer. 2009. Phenomics: the systematic study of phenotypes on a genome-wide scale. Neuroscience 164:30–42.

Blonder, B. 2018. Hypervolume concepts in niche- and trait-based ecology. Ecography 41:1441–1455.

Breiman, L. 2001. Random Forests. Machine learning 45:5–32.

Brombacher, A., P. A. Wilson, I. Bailey, and T. H. G. Ezard. 2017. The Breakdown of Static and Evolutionary Allometries during Climatic Upheaval. The American naturalist 190:350–362.

Bruijning, M., M. D. Visser, C. A. Hallmann, and E. Jongejans. 2018. trackdem : Automated particle tracking to obtain population counts and size distributions from videos in r. Methods in ecology and evolution / British Ecological Society 9:965–973.

Buetti-Dinh, A., V. Galli, S. Bellenberg, O. Ilie, M. Herold, S. Christel, M. Boretska, I. V. Pivkin, P. Wilmes, W. Sand, M. Vera, and M. Dopson. 2019. Deep neural networks outperform human expert's capacity in characterizing bioleaching bacterial biofilm composition. Biotechnology reports (Amsterdam, Netherlands) 22:e00321.

Canny, J. 1986. A Computational Approach to Edge Detection. IEEE transactions on pattern analysis and machine intelligence PAMI-8:679–698.

Cheng, K. C., X. Xin, D. P. Clark, and P. La Riviere. 2011. Whole-animal imaging, gene function, and the Zebrafish Phenome Project. Current opinion in genetics & development 21:620–629.

Chesson, P. 2000. Mechanisms of Maintenance of Species Diversity. Annual review of

649 ecology and systematics 31:343–366.
650 Chevin, L.-M., R. Lande, and G. M. Mace. 2010. Adaptation, plasticity, and extinction in a
651 changing environment: towards a predictive theory. PLoS biology 8:e1000357.
652 Church, G. M., and W. Gilbert. 1984. Genomic sequencing. Proceedings of the National
653 Academy of Sciences of the United States of America 81:1991–1995.
654 Church, S. H., S. Donoughe, B. A. S. de Medeiros, and C. G. Extavour. 2019a, July. Insect
655 egg size and shape evolve with ecology but not developmental rate.
656 Church, S. H., S. Donoughe, B. A. S. de Medeiros, and C. G. Extavour. 2019b. A dataset of
657 egg size and shape from more than 6,700 insect species. Scientific data 6:104.
658 Cortes, C., and V. Vapnik. 1995. Support-Vector Networks. Machine learning 20:273–297.
659 Dalal, N., and B. Triggs. 2005. Histograms of oriented gradients for human detection. Pages
660 886–893 vol. 1 2005 IEEE Computer Society Conference on Computer Vision and
661 Pattern Recognition (CVPR'05).
662 Dietterich, T. G. 2000. Ensemble Methods in Machine Learning. Pages 1–15 Proceedings of
663 the First International Workshop on Multiple Classifier Systems. Springer-Verlag, Berlin,
664 Heidelberg.
665 Di, Z., M. J. D. Klop, V.-M. Rogkoti, S. E. Le Dévédec, B. van de Water, F. J. Verbeek, L. S.
666 Price, and J. H. N. Meerman. 2014. Ultra high content image analysis and phenotype
667 profiling of 3D cultured micro-tissues. PloS one 9:e109688.
668 Duda, R. O., and P. E. Hart. 1972. Use of the Hough transformation to detect lines and
669 curves in pictures. Communications of the ACM 15:11–15.
670 Feder, M. E., and T. Mitchell-Olds. 2003. Evolutionary and ecological functional genomics.
671 Nature reviews. Genetics 4:651–657.
672 Freimer, N., and C. Sabatti. 2003. The human phenome project. Nature genetics 34:15–21.
673 French, S., B. E. Coutts, and E. D. Brown. 2018. Open-Source High-Throughput Phenomics
674 of Bacterial Promoter-Reporter Strains. Cell systems 7:339–346.e3.
675 Friedman, J. H. 2000. Greedy Function Approximation: A Gradient Boosting Machine.
676 Annals of Statistics.

677 Furbank, R. T., and M. Tester. 2011. Phenomics--technologies to relieve the phenotyping
678 bottleneck. Trends in plant science 16:635–644.
679 Gehan, M. A., N. Fahlgren, A. Abbasi, J. C. Berry, S. T. Callen, L. Chavez, A. N. Doust, M. J.
680 Feldman, K. B. Gilbert, J. G. Hodge, J. S. Hoyer, A. Lin, S. Liu, C. Lizárraga, A.
681 Lorence, M. Miller, E. Platon, M. Tessman, and T. Sax. 2017. PlantCV v2: Image
682 analysis software for high-throughput plant phenotyping. PeerJ 5:e4088.
683 Geoffrey Hinton, Terrence J. Sejnowski. 1999. Unsupervised Learning: Foundations of
684 Neural Computation. MIT Press.
685 Gerum, R. C., S. Richter, B. Fabry, and D. P. Zitterbart. 2017. ClickPoints: an expandable
686 toolbox for scientific image annotation and analysis. Methods in ecology and evolution /
687 British Ecological Society 8:750–756.
688 Grabowski, M., and A. Porto. 2017. How many more? Sample size determination in studies
689 of morphological integration and evolvability. Methods in ecology and evolution / British
690 Ecological Society 8:592–603.
691 Hakim, A., Y. Mor, I. A. Toker, A. Levine, M. Neuhof, Y. Markovitz, and O. Rechavi. 2018.
692 WorMachine: machine learning-based phenotypic analysis tool for worms. BMC biology
693 16:8.
694 Heileman, G. L., and H. R. Myler. 1989. Theoretical and experimental aspects of supervised
695 learning in artificial neural networks. phd, University of Central Florida, USA.
696 Hense, B. A., P. Gais, U. Jutting, H. Scherb, and K. Rodenacker. 2008. Use of fluorescence
697 information for automated phytoplankton investigation by image analysis.
698 Hoffmann, J., S. Donoughe, K. Li, M. K. Salcedo, and C. H. Rycroft. 2018. A simple
699 developmental model recapitulates complex insect wing venation patterns. Proceedings
700 of the National Academy of Sciences of the United States of America 115:9905–9910.
701 Houle, D., G. H. Bolstad, K. van der Linde, and T. F. Hansen. 2017. Mutation predicts 40
702 million years of fly wing evolution. Nature 548:447–450.
703 Houle, D., D. R. Govindaraju, and S. Omholt. 2010. Phenomics: the next challenge. Nature
704 Reviews. Genetics 11:855–866.

Houle, D., J. Mezey, P. Galpern, and A. Carter. 2003. Automated measurement of Drosophila wings. BMC evolutionary biology 3:25.

Houle, D., C. Pélabon, G. P. Wagner, and T. F. Hansen. 2011. Measurement and Meaning in Biology.

Hsiang, A. Y., K. Nelson, L. E. Elder, E. C. Sibert, S. S. Kahanamoku, J. E. Burke, A. Kelly, Y. Liu, and P. M. Hull. 2018. AutoMorph : Accelerating morphometrics with automated 2D and 3D image processing and shape extraction. Methods in ecology and evolution / British Ecological Society 9:605–612.

Ishikawa, A., N. Kabeya, K. Ikeya, R. Kakioka, J. N. Cech, N. Osada, M. C. Leal, J. Inoue, M. Kume, A. Toyoda, A. Tezuka, A. J. Nagano, Y. Y. Yamasaki, Y. Suzuki, T. Kokita, H. Takahashi, K. Lucek, D. Marques, Y. Takehana, K. Naruse, S. Mori, O. Monroig, N. Ladd, C. J. Schubert, B. Matthews, C. L. Peichel, O. Seehausen, G. Yoshizaki, and J. Kitano. 2019. A key metabolic gene for recurrent freshwater colonization and radiation in fishes. Science 364:886–889.

Kanopoulos, N., N. Vasanthavada, and R. L. Baker. 1988. Design of an image edge detection filter using the Sobel operator. IEEE journal of solid-state circuits 23:358–367.

Kell, D. B., and S. G. Oliver. 2004. Here is the evidence, now what is the hypothesis? The complementary roles of inductive and hypothesis-driven science in the post-genomic era. BioEssays: news and reviews in molecular, cellular and developmental biology 26:99–105.

Kingsolver, J. G., and R. B. Huey. 2008. Size, temperature, and fitness: three rules. Evolutionary ecology research 10:251–268.

Kühl, H. S., and T. Burghardt. 2013. Animal biometrics: quantifying and detecting phenotypic appearance. Trends in ecology & evolution 28:432–441.

Lamichhaney, S., D. C. Card, P. Grayson, J. F. R. Tonini, G. A. Bravo, K. Näpflin, F. Termignoni-Garcia, C. Torres, F. Burbrink, J. A. Clarke, T. B. Sackton, and S. V. Edwards. 2019. Integrating natural history collections and comparative genomics to study the genetic architecture of convergent evolution. Philosophical transactions of the

733 Royal Society of London. Series B, Biological sciences 374:20180248.
734 Lande, R. 2009. Adaptation to an extraordinary environment by evolution of phenotypic
735 plasticity and genetic assimilation. Journal of evolutionary biology 22:1435–1446.
736 Lande, R., and S. J. Arnold. 1983. The Measurement of Selection on Correlated Characters.
737 Evolution 37:1210–1226.
738 Laughlin, D. C., J. R. Gremer, P. B. Adler, R. M. Mitchell, and M. M. Moore. 2020. The Net
739 Effect of Functional Traits on Fitness. Trends in ecology & evolution 35:1037–1047.
740 LeCun, Y., Y. Bengio, and G. Hinton. 2015. Deep learning. Nature 521:436–444.
741 Leighton, G., P. S. Hugo, A. Roulin, and A. Amar. 2016. Just Google it: assessing the use of
742 Google Images to describe geographical variation in visible traits of organisms 7.
743 Le, V.-L., M. Beurton-Aimar, A. Zemmari, A. Marie, and N. Parisey. 2020. Automated
744 landmarking for insects morphometric analysis using deep neural networks. Ecological
745 informatics 60:101175.
746 Liow, L. H., E. Di Martino, M. Krzeminska, M. Ramsfjell, S. Rust, P. D. Taylor, and K. L.
747 Voje. 2017. Relative size predicts competitive outcome through 2 million years. Ecology
748 letters 20:981–988.
749 Liu, F., A. Wollstein, P. G. Hysi, G. A. Ankra-Badu, T. D. Spector, D. Park, G. Zhu, M.
750 Larsson, D. L. Duffy, G. W. Montgomery, D. A. Mackey, S. Walsh, O. Lao, A. Hofman,
751 F. Rivadeneira, J. R. Vingerling, A. G. Uitterlinden, N. G. Martin, C. J. Hammond, and
752 M. Kayser. 2010. Digital Quantification of Human Eye Color Highlights Genetic
753 Association of Three New Loci.
754 Lloyd, S. 1982. Least squares quantization in PCM. IEEE transactions on information theory
755 / Professional Technical Group on Information Theory 28:129–137.
756 Lowe, D. G. 1999. Object recognition from local scale-invariant features. Pages 1150–1157
757 vol.2 Proceedings of the Seventh IEEE International Conference on Computer Vision.
758 Lowe, D. G. 2004. Distinctive image features from scale-invariant keypoints. International
759 journal of computer vision.
760 Lürig, M. D. 2018. phenopype - a phenotyping pipeline for python.

761 Maeda, T., J. Iwasawa, H. Kotani, N. Sakata, M. Kawada, T. Horinouchi, A. Sakai, K.
762 Tanabe, and C. Furusawa. 2020. High-throughput laboratory evolution reveals
763 evolutionary constraints in Escherichia coli. Nature communications 11:5970.
764 McPeek, M. A., L. Shen, J. Z. Torrey, and H. Farid. 2008. The Tempo and Mode of Three-
765 Dimensional Morphological Evolution in Male Reproductive Structures.
766 McQuin, C., A. Goodman, V. Chernyshev, L. Kamentsky, B. A. Cimini, K. W. Karhohs, M.
767 Doan, L. Ding, S. M. Rafelski, D. Thirstrup, W. Wiegraebe, S. Singh, T. Becker, J. C.
768 Caicedo, and A. E. Carpenter. 2018. CellProfiler 3.0: Next-generation image processing
769 for biology. PLoS biology 16:e2005970.
770 Minsky, M. 1961. Steps toward Artificial Intelligence. Proceedings of the IRE 49:8–30.
771 Minsky, M, Papert, and S. 1969. Perceptrons.
772 Mitchell, T. M. 1997. Machine learning. 1997. Burr Ridge, IL: McGraw Hill 45:870–877.
773 Morel-Journel, T., V. Thuillier, F. Pennekamp, E. Laurent, D. Legrand, A. S. Chaine, and N.
774 Schtickzelle. 2020. A multidimensional approach to the expression of phenotypic
775 plasticity. Functional ecology.
776 O'Mahony, N., S. Campbell, A. Carvalho, S. Harapanahalli, G. V. Hernandez, L. Krpalkova,
777 D. Riordan, and J. Walsh. 2020. Deep Learning vs. Traditional Computer Vision.
778 Orgogozo, V., B. Morizot, and A. Martin. 2015. The differential view of genotype-phenotype
779 relationships. Frontiers in genetics 6:179.
780 Pfennig, D. W., M. A. Wund, E. C. Snell-Rood, T. Cruickshank, C. D. Schlichting, and A. P.
781 Moczek. 2010. Phenotypic plasticity's impacts on diversification and speciation. Trends
782 in ecology & evolution 25:459–467.
783 Phillips, C. P., and S. J. Arnold. 1999a. Hierarchical comparison of genetic variance
784 covariance matrices. I. Using the Flury hierarchy. Evolution; international journal of
785 organic evolution 53:1506–1515.
786 Phillips, P. C., and S. J. Arnold. 1999b. Hierarchical Comparison of Genetic Variance-
787 Covariance Matrices. I. Using the Flury Hierarchy. Evolution; international journal of
788 organic evolution 53:1506–1515.

789 Pitchers, W., J. Nye, E. J. Márquez, A. Kowalski, I. Dworkin, and D. Houle. 2019. A
790 Multivariate Genome-Wide Association Study of Wing Shape in Drosophila
791 melanogaster. Genetics 211:1429–1447.
792 Pointer, M. R., and G. G. Attridge. 1998. The number of discernible colours. Color Research
793 & Application: Endorsed by Inter-Society Color Council, The Colour Group (Great
794 Britain), Canadian Society for Color, Color Science Association of Japan, Dutch Society
795 for the Study of Color, The Swedish Colour Centre Foundation, Colour Society of
796 Australia, Centre Français de la Couleur 23:52–54.
797 Porto, A., and K. Lysne Voje. 2020. ML-morph: A fast, accurate and general approach for
798 automated detection and landmarking of biological structures in images. Methods in
799 ecology and evolution / British Ecological Society.
800 Porto, A., S. M. Rolfe, and A. Murat Maga. 2020, September 19. ALPACA: a fast and
801 accurate approach for automated landmarking of three-dimensional biological
802 structures.
803 Porto, A., and K. L. Voje. 2020. ML-morph: A fast, accurate and general approach for
804 automated detection and landmarking of biological structures in images. Methods in
805 ecology and evolution / British Ecological Society 11:500–512.
806 Reynolds, D. A., and R. C. Rose. 1995. Robust text-independent speaker identification using
807 Gaussian mixture speaker models. IEEE transactions on audio, speech, and language
808 processing 3:72–83.
809 Ringnér, M. 2008. What is principal component analysis? Nature biotechnology 26:303–304.
810 Roberts, L. G. 1963. Machine perception of three-dimensional solids. Massachusetts
811 Institute of Technology.
812 Roeder, A. H. K., A. Cunha, M. C. Burl, and E. M. Meyerowitz. 2012. A computational image
813 analysis glossary for biologists. Development 139:3071–3080.
814 Rosenblatt, F. 1958. The perceptron: a probabilistic model for information storage and
815 organization in the brain. Psychological review 65:386–408.
816 Rumelhart, D. E., and J. L. McClelland. 1987. Learning Internal Representations by Error

817 Propagation. Pages 318–362 Parallel Distributed Processing: Explorations in the
818 Microstructure of Cognition: Foundations. MIT Press.
819 Salcedo, M. K., J. Hoffmann, S. Donoughe, and L. Mahadevan. 2019. Computational
820 analysis of size, shape and structure of insect wings. Biology open 8.
821 Saltz, J. B., F. C. Hessel, and M. W. Kelly. 2017. Trait Correlations in the Genomics Era.
822 Trends in ecology & evolution.
823 Schindelin, J., I. Arganda-Carreras, E. Frise, V. Kaynig, M. Longair, T. Pietzsch, S.
824 Preibisch, C. Rueden, S. Saalfeld, B. Schmid, J.-Y. Tinevez, D. J. White, V.
825 Hartenstein, K. Eliceiri, P. Tomancak, and A. Cardona. 2012. Fiji: an open-source
826 platform for biological-image analysis. Nature methods 9:676–682.
827 Schluter, D. 1996. Adaptive radiation along genetic lines of least resistance. Evolution;
828 international journal of organic evolution 50:1766–1774.
829 Seehausen, O., R. K. Butlin, I. Keller, C. E. Wagner, J. W. Boughman, P. A. Hohenlohe, C.
830 L. Peichel, G. P. Saetre, C. Bank, A. Brannstrom, A. Brelsford, C. S. Clarkson, F.
831 Eroukhmanoff, J. L. Feder, M. C. Fischer, A. D. Foote, P. Franchini, C. D. Jiggins, F. C.
832 Jones, A. K. Lindholm, K. Lucek, M. E. Maan, D. A. Marques, S. H. Martin, B.
833 Matthews, J. I. Meier, M. Most, M. W. Nachman, E. Nonaka, D. J. Rennison, J.
834 Schwarzer, E. T. Watson, A. M. Westram, and A. Widmer. 2014. Genomics and the
835 origin of species. Nature reviews. Genetics 15:176–192.
836 Sinervo, B., and E. Svensson. 2002. Correlational selection and the evolution of genomic
837 architecture. Heredity 89:329–338.
838 Soulé, M. 1967. PHENETICS OF NATURAL POPULATIONS I. PHENETIC
839 RELATIONSHIPS OF INSULAR POPULATIONS OF THE SIDE-BLOTCHED LIZARD.
840 Evolution; international journal of organic evolution 21:584–591.
841 Svensson, E. I., S. J. Arnold, R. Bürger, K. Csilléry, J. Draghi, J. M. Henshaw, A. G. Jones,
842 S. De Lisle, D. A. Marques, K. McGuigan, M. N. Simon, and A. Runemark. 2021.
843 Correlational selection in the age of genomics. Nature Ecology and Evolution.
844 Svensson, E. I., M. Gomez-Llano, and J. Waller. 2020, June 12. Natural and sexual

selection on phenotypic plasticity favour thermal canalization.

Svensson, E. I., and J. T. Waller. 2013. Ecology and sexual selection: evolution of wing pigmentation in calopterygid damselflies in relation to latitude, sexual dimorphism, and speciation. The American naturalist 182:E174–95.

Tattersall, G. J., D. V. Andrade, and A. S. Abe. 2009. Heat exchange from the toucan bill reveals a controllable vascular thermal radiator. Science 325:468–470.

Tattersall, G. J., and V. Cadena. 2010. Insights into animal temperature adaptations revealed through thermal imaging. The Imaging Science Journal 58:261–268.

Tsubaki, Y., Y. Samejima, and M. T. Siva-Jothy. 2010. Damselfly females prefer hot males: higher courtship success in males in sunspots. Behavioral ecology and sociobiology 64:1547–1554.

Tsuboi, M., B. T. Kopperud, C. Syrowatka, M. Grabowski, K. L. Voje, C. Pélabon, and T. F. Hansen. 2020. Measuring Complex Morphological Traits with 3D Photogrammetry: A Case Study with Deer Antlers. Evolutionary biology 47:175–186.

Turk, M., and A. Pentland. 1991. Eigenfaces for recognition. Journal of cognitive neuroscience 3:71–86.

Ubbens, J. R., and I. Stavness. 2017. Deep Plant Phenomics: A Deep Learning Platform for Complex Plant Phenotyping Tasks. Frontiers in plant science 8:1190.

Visscher, P. M., and J. Yang. 2016. A plethora of pleiotropy across complex traits.

Walsh, B. 2007, January. Escape from flatland.

Weinstein, B. G. 2015. MotionMeerkat: integrating motion video detection and ecological monitoring. Methods in ecology and evolution / British Ecological Society 6:357–362.

Wolak, M. E., D. J. Fairbairn, and Y. R. Paulsen. 2012. Guidelines for estimating repeatability. Methods in ecology and evolution / British Ecological Society 3:129–137.

Wolman, A. G. 2006. Measurement and meaningfulness in conservation science. Conservation biology: the journal of the Society for Conservation Biology 20:1626–1634.

Zackrisson, M., J. Hallin, L.-G. Ottosson, P. Dahl, E. Fernandez-Parada, E. Ländström, L.

873 Fernandez-Ricaud, P. Kaferle, A. Skyman, S. Stenberg, S. Omholt, U. Petrovič, J.
874 Warringer, and A. Blomberg. 2016. Scan-o-matic: High-Resolution Microbial Phenomics
875 at a Massive Scale. G3 6:3003–3014.
876 Zheng, J., J. L. Payne, and A. Wagner. 2019. Cryptic genetic variation accelerates evolution
877 by opening access to diverse adaptive peaks. Science 365:347–353.

878

Figures

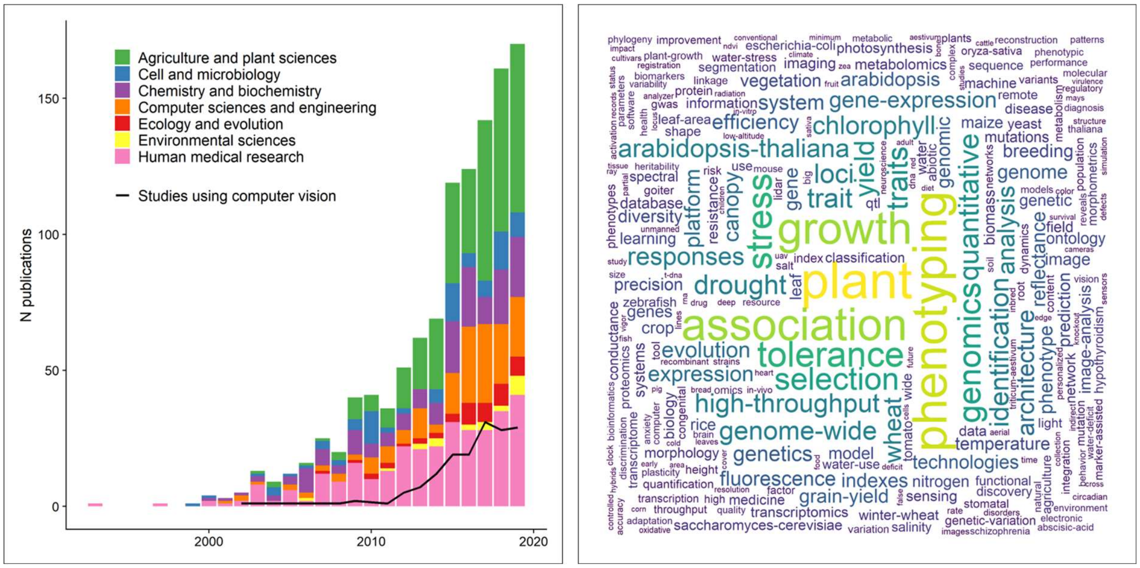

Figure 1 - The current state of phenomics research. Left panel: a literature survey using the search key “phenomic* on in a Web of Science topic search (title, keywords, abstract) resulted in 1323 papers (on 23/10/2020). Here we show only papers published between 1990 and before 2020 (1125 papers) for better visual inference. Traditionally phenomics approaches are used in agricultural sciences and crop research to maximize yield, as well as in human medicine to study drug responses and disorder phenotypes. The black line denotes the studies that used computer vision or some sort of image analysis (acquired with a topic search using the strings "computer vision", "image analysis", "image based", "image processing"), indicating that only a small subset of the studies uses image analysis. Right panel: a word cloud that was constructed using the 500 most used keywords from the papers presented in the left panel.

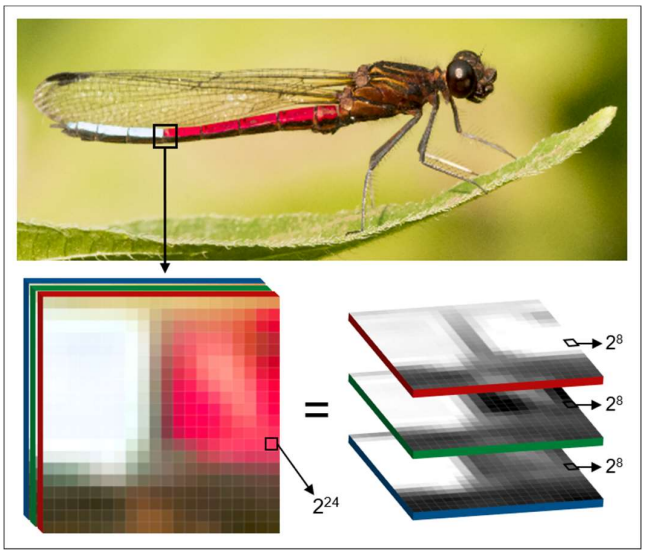

Figure 2 - The structure of digital images. Two-dimensional raster images, as produced by most commercially available cameras, are composed of three color channels red, green, blue = RGB), each of which by itself is a grayscale image. The industrial standard for color representation on the pixel level is 24 bit (2^{24} = 16 777 216 possible color variations per pixel), which is achieved through additive mixing of each of the 8 bit channels ($2^8 + 2^8 + 2^8$). This enormous range of color intensities among several million pixels is a potentially very high-resolution representation of organismal traits, or the organism as a whole. Therefore, digital images are a useful medium for phenomics research, as they offer an inexpensive, memory efficient and standardizable way to capture, store and analyze complex phenotypes. The photograph shows a blue-tip jewel damselfly (*Chlorocypha curta*) in Cameroon (Africa) - image by Erik Svensson.

911

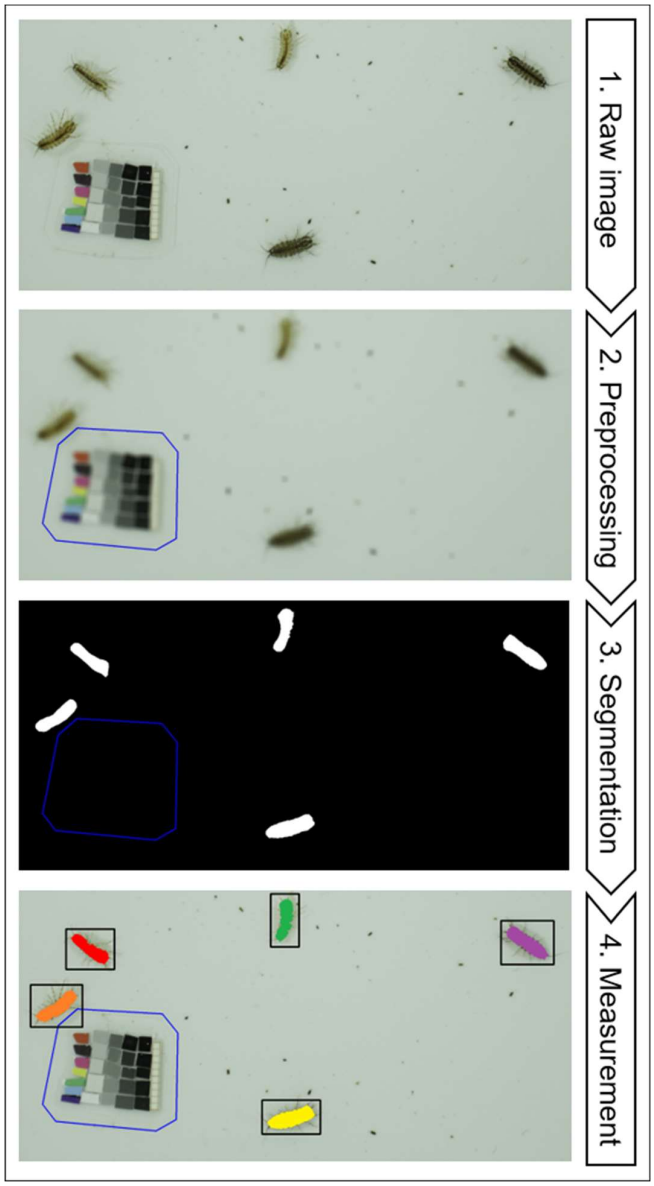

912

Figure 3 - A typical computer vision workflow using signal processing. 1) Raw image - The goal is to detect, count and measure freshwater isopods (*Asellus aquaticus*, image by Moritz Lürig) from the raw image that was taken under controlled laboratory conditions. 2) Preprocessing - The operating principle of most signal processing workflows is that the objects of interest are made to contrast strongly from all other pixels, meaning that images should have a high *signal-to-noise ratio* (SNR. In this specific case a high SNR is already present, because the isopods are much darker than the tray they are sitting on and much larger than the fecal pellets and other detritus around them. To further increase the SNR, gaussian blur blends pixels in a given neighborhood (=*kernel* size), which effectively removes the smaller dark objects. The reference card gets excluded manually and can be used to convert pixels to millimeters and to correct the color space. 3) Segmentation - Using a thresholding algorithm all connected pixels that are above a specific grayscale value and larger than a specified area are designated foreground (white) and all pixels become background (black). The output from this step is referred to as a “binary mask”. 4) Measurement - Now the white pixels from the binary mask can be used to locate the areas of interest in the raw image and to extract information from them. Discrimination between multiple instances of the same class is referred to as instance segmentation.

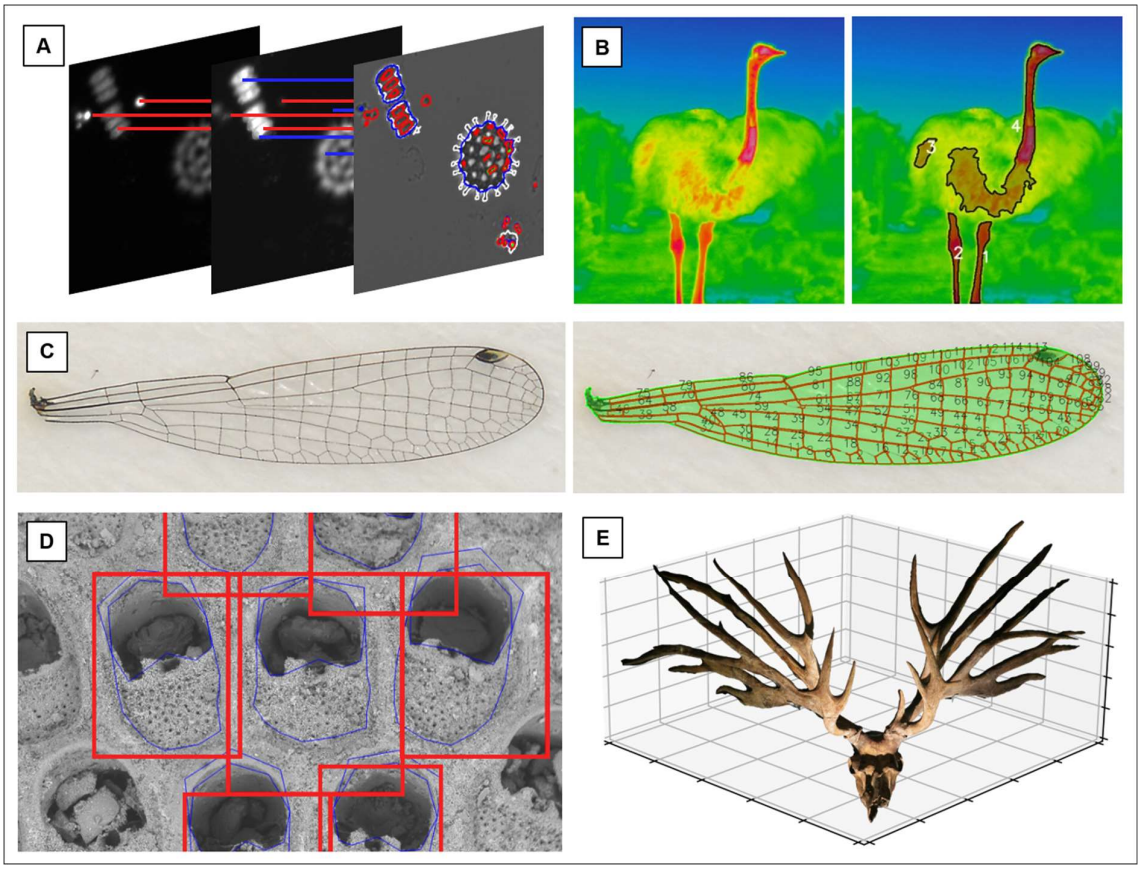

Figure 4 - Different types of high dimensional phenotypic data that are collected using a fully or semi-automatic computer vision approach. A) Morphology and fluorescence traits of phytoplankton communities are represented through a combination of shape features (e.g. circularity, perimeter length, area) and texture features (e.g. blob intensity and distribution within the cell) from images showing fluorescence intensity (pictograms on the left; images by Irene Gallego). B) In ostriches (*Struthio camelus*), surface temperatures of bare body parts without feathers (necks and legs) are detected using signal processing (image by Erik Svensson). C) Signal processing approach that captures individual domains of a damselfly wing via thresholding (image by Masahito Tsuboi). D) Ensemble-based approach to shape prediction of individual zooids within a bryozoan colony (image by Arthur Porto) E) 3D image of the skull of extinct deer *Eucladoceros dicranios* from which we can measure linear, area, and volumetric measurements of antler features (image by Masahito Tsuboi).

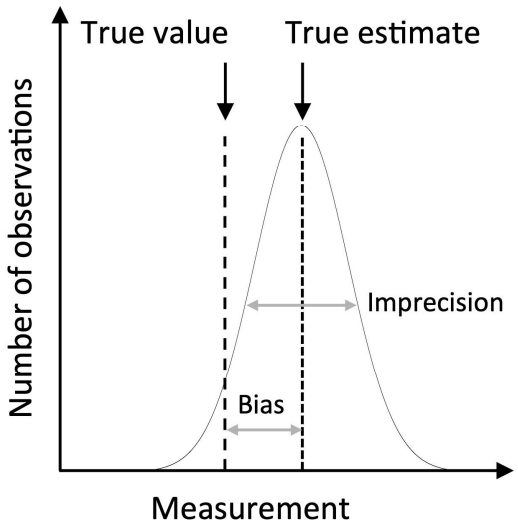

Figure 5 - Schematic illustration of bias and imprecision. X-axis represents phenotypic values and Y-axis represents number of observations. The gaussian curve shows the distribution of repeated measurements of the same specimen. Dashed line is the true estimate, and the variance of measurements around the true estimate is the imprecision. The true value may deviate systematically from the true estimate (long-dashed line). The difference between true estimate and true value is the bias.

A

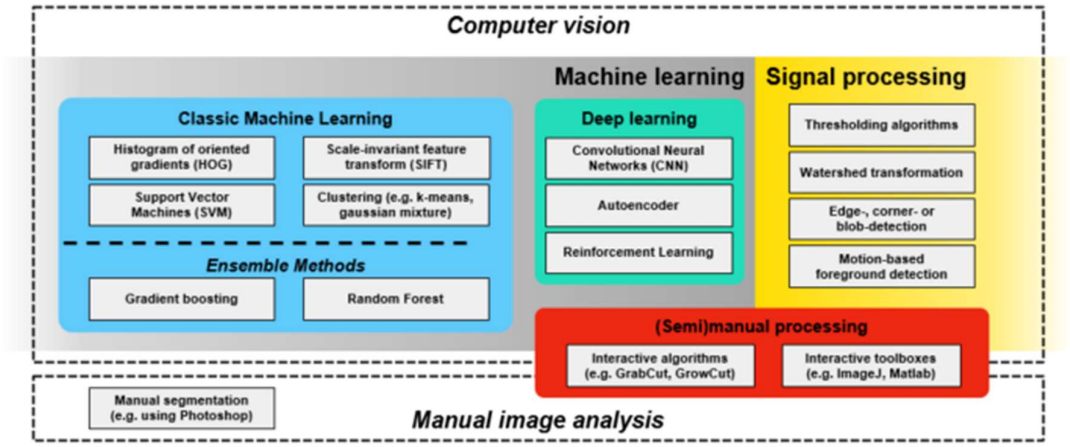
Computer vision
Machine learning
Signal processing
Classic Machine Learning
Histogram of oriented gradients (HOG)
Scale-invariant feature transform (SIFT)
Support Vector Machines (SVM)
Clustering (e.g. k-means, gaussian mixture)
Ensemble Methods
Gradient boosting
Random Forest
Deep learning
Convolutional Neural Networks (CNN)
Autoencoder
Reinforcement Learning
Thresholding algorithms
Watershed transformation
Edge-, corner- or blob-detection
Motion-based foreground detection
(Semi)manual processing
Interactive algorithms (e.g. GrabCut, GrowCut)
Interactive toolboxes (e.g. ImageJ, Matlab)
Manual segmentation (e.g. using Photoshop)
Manual image analysis

B

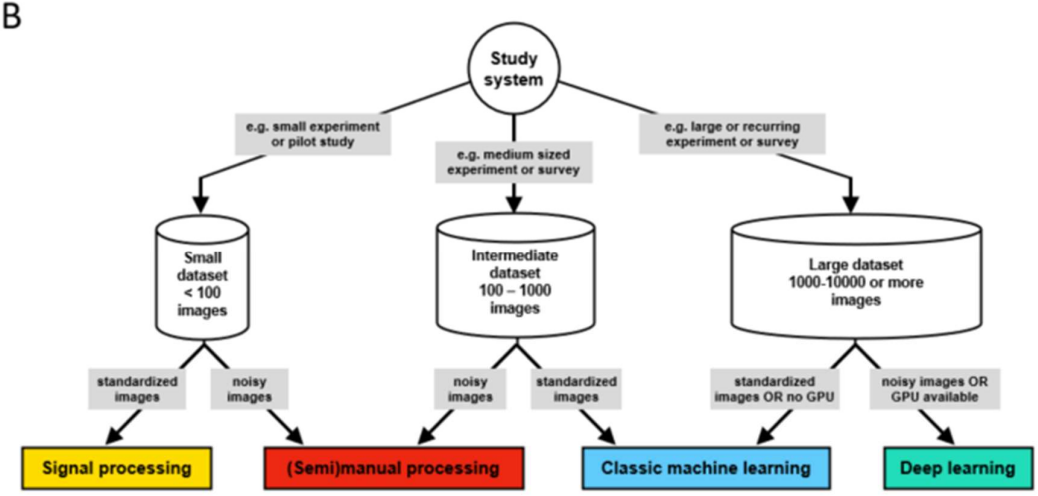
Study system
e.g. small experiment or pilot study
e.g. medium sized experiment or survey
e.g. large or recurring experiment or survey
Small dataset < 100 images
Intermediate dataset 100 – 1000 images
Large dataset 1000-10000 or more images
standardized images
noisy images
noisy images
standardized images
standardized images OR no GPU
noisy images OR GPU available
Signal processing
(Semi)manual processing
Classic machine learning
Deep learning

950

Figure 6 - Computer vision (CV) methods overview - which is the right one for my data? A) CV is a field at the intersection of machine learning and signal processing which is concerned with the automatic and semiautomatic extraction of information from digital images. B) Decision tree for CV methods: begin by considering the size of a given image dataset, whether it is complete, e.g. after an experiment, or whether there will be continued future additions, e.g. as part of a long-term experiment or field survey. The next assessment should be about the signal-to-noise ratio (SNR) in your images: images taken in the laboratory typically have a high degree of standardization and thus a very high SNR, which makes them suitable for a signal processing approach. In contrast to signal processing, where segmentation results are immediately available, all machine learning image analysis pipelines include iterative training and validation phases, followed by a final testing phase. Such a workflow can be complex to initiate but pays off in the long run by providing segmentation results that become increasingly robust if more training data is supplied over time. Deep learning algorithms require large training datasets (several 1000s to 10000s) but are less sensitive to noise and idiosyncrasies of the foreground. Thus, for large and continuously growing data sets, or for recurring image analysis tasks, deep learning has become the standard approach for segmentation.

Tables

Table 1 - Classes of phenotypic data. Depending on the research question, scientists define their phenotypes of interest using specific or abstract, low or high dimensional traits (see section *Measurement theory*). The human eye excels at rapidly recognizing externally visible phenotypes (e.g. benthic vs. limnetic morphotypes of fish), but has difficulties discerning what constitutes such phenotypes. Computer vision offers an objective way to collect any data type with high efficiency and reproducibility. For instance, by breaking down low dimensional traits (e.g. red vs. blue phenotype) into continuous low or high dimensional metrics (e.g. degree of red- or blueness), the decision of what constitutes a phenotype becomes more reproducible.

Trait type	**Low dimensional**	**High dimensional**
Specific / directly measurable	Size, discrete color (“red phenotype” vs. “blue phenotype”) and morphotype scoring (e.g. benthic vs limnetic)	Shape coordinates, texture maps, landmarks
Abstract / derived	Shape (e.g. circularity, area) and texture features (e.g. mean, SD, uniformity), moments, principal components, hypervolumes	Matrices, activation maps

Table 2 - Select examples of recent open-source computer vision libraries with a biology-context. Although typically first developed for a particular study system or organism (e.g. PlantCV or WorMachine), most CV applications apply techniques that are generally applicable to any type of phenotypic data contained in digital images.

Name	Year	Reference	Repository	Purpose	Application type	Description	Techniques
AutoMorph	2018	(Hsiang et al. 2018)	https://github.com/HullLab/AutoMorph	object detection and feature extraction	Python package	High throughput segmentation	Signal processing
ClickPoints	2017	(Gerum et al. 2017)	https://github.com/fabrylab/clickpoints	labelling, label evaluation	Python package	Interactive labelling tool	signal processing
DeepMerkat	2018	(Weinstein 2015)	https://github.com/bw4sz/DeepMeerkat	object detection, classification	Python	Background subtraction and image classification for stationary cameras in ecological videos	Signal processing, deep learning
EB-Net	2020	(Le et al. 2020)	https://github.com/linhlevandlu/CNN_Beetles_Landmarks	keypoint and feature detection	Python	Insect morphometrics	deep learning
ImageJ	2012	(Schindelin et al. 2012)	https://fiji.sc/; https://imagej.nih.gov/ij/download.html	multi-purpose	standalone	Comprehensive, multi-purpose image processing library	manual processing, signal processing, classic machine learning, feature extraction
ML-morph	2020	(Porto and Lysne Voje 2020)	https://github.com/agporto/ml-morph	landmark detection; geometric morphometrics	Python package	High throughput morphometrics	Classic machine learning, ensemble Methods
MotionMeerkat	2015	(Weinstein 2015)	https://github.com/bw4sz/DeepMeerkat	motion tracking	Python package/ standalone	Deep learning driven motion detection	Signal processing, deep learning
Phenopype	2020	(Lürig 2018)	https://github.com/mluerig/phenopype	object detection, feature extraction, motion tracking	Python package	Computer vision library with high throughput workflows	signal processing

PlantCV	2017	(Gehan et al. 2017)	https://github.com/danforthcenter/plantcv	object detection and feature extraction; spectral analysis	Python package	Plant phenotyping library	signal processing, classic machine learning
Scan-o-matic	2016	(Zackrisson et al. 2016)	https://github.com/Scan-o-Matic/scanomatic	object detection and feature extraction	Python package	Microbial phenotyping platform	Signal processing
Trackdem	2017	(Bruijning et al. 2018)	https://github.com/marjoleinbruijning/trackdem	motion tracking and blob counting	R package	Behavioral analysis pipeline	Signal processing
WingMachine	2003	(Houle et al. 2003)	https://www.bio.fsu.edu/~dhoule/Software/	keypoint and feature detection	standalone	Drosophila wing morphometrics	Signal processing, feature extraction
WorMachine	2018	(Hakim et al. 2018)	https://github.com/adamhak/WorMachineClient	object detection and feature extraction	Matlab	Integrated image processing and feature extraction	Signal processing, classic machine learning; deep learning

986

987

988

989

990 # Boxes

991 Box 1 - Glossary of terms relevant for computer vision and machine learning in ecology and
992 evolution used in this review. Terms in this list are printed in *italic* when first mentioned in the
993 main text.

994

bit depth	number of values a pixel can take (e.g. 8 bit = 2^8 = 256 values)
computer vision	technical domain at the intersection of signal processing, machine learning, robotics and other scientific areas that is concerned with the automated extraction of information from digital images and videos.
convolution	mathematical operation by which information contained in images are abstracted. Each convolutional layer produces a feature map, which is passed on to the next layer.
deep learning	machine learning methods based on neural networks. supervised learning = algorithm learns input features from input-output pairs (e.g. labelled images). unsupervised = algorithm looks for undetected patterns (e.g. images without labelling)
feature	a measurable property or pattern. can be specific (e.g. edges, corners, points) or abstract (e.g. convolution via kernels), and combined to vectors and matrices (feature maps)
feature detection	methods for making pixel-level or pixel-neighborhood decisions on whether parts of an image are a feature or not
foreground	all pixels of interest in a given image, whereas the background constitutes all other pixels. the central step in computer vision is the segmentation of all pixels into foreground and background
hidden layer	a connected processing step in neural networks during which information is received, processed (e.g. convolved), and passed on to the next layer
kernel	a small mask or matrix to perform operations on images, for example, blurring, sharpening or edge detection. the kernel operation is performed pixel wise, sliding across the entire image.
labelling	typically manual markup of areas of interest in an image by drawing bounding boxes or polygons around the contour. can be multiple objects and multiple classes of objects per image. can also refer to assigning whole images to a class (e.g. relevant for species identification)
machine learning	subset of artificial intelligence: the study and implementation of computer algorithms that improve automatically through experience. (Mitchell 1997)
measurement theory	concerns the relationship between measurements and nature so that inferences from measurements reflect the underlying reality intended to be represent (Houle et al. 2011).
neural network	deep learning algorithms that use multi layered ("deep") abstractions of information to extract higher level features from input via convolution

object detection	methods for determining whether a pixel region constitutes an object that belongs to the foreground or not, based on its features
phenome	the phenotype as a whole (sensu Soulé 1967)
phenomics	the acquisition of high-dimensional phenotypic data on an organism-wide scale
phenotype	a single trait or a specific set of traits that are part of the phenome
pixel	short for picture element; the smallest accessible unit of a digital raster image. Pixels have finite values (=intensities), e.g. 256 in an 8-bit grayscale image.
segmentation	the classification of all pixels in an image into foreground and background, either manually by labelling the area of interest, or automatically, by means of signal processing or machine learning algorithms. semantic segmentation = all pixels of a class, instance segmentation = all instances of a class
signal processing	technically correct: digital image processing (not to be confused with image analysis or image editing). subfield of engineering that is concerned with the filtering or modification of digital images by means of algorithms and filter matrices (kernels),
signal-to-noise ratio (SNR)	describes the level of the pixels containing the desired signal (i.e. the phenotypic information) to all other pixels. Lab images typically have a high SNR, field images a low SNR.
threshold algorithm	pixel-intensity based segmentation of images, e.g. based on individual pixel intensity (binary thresholding) or their intensity with respect to their neighborhood (adaptive thresholding). creates a binary mask which contains only black or white pixels
training data	representative image dataset to train a machine learning algorithm. can be created manually by labelling images, or semi-automatic by using signal processing for segmentation. can contain single or multiple classes
watershed algorithm	the segmentation of images by treating the pixels as a topographic map of basins, where bright pixels have high elevation and dark pixels have low elevation.

Box 2 - An overview of the main deep learning architectures and approaches.

Families of network topologies

A. **Deep convolutional network** - A large and common family of neural networks composed an input layer, an output layer and multiple hidden layers. These networks feature convolution kernels that process input data and pooling layers that simplify the information processed through the convolutional kernels. For certain tasks, the input can be a window of the image, rather than the entire image.
B. **Deconvolutional Network** - A smaller family of neural networks that perform the reverse process when compared to convolutional networks. It starts with the processed data (i.e.., the output of the convolutional network) and it aims to separate what has been convoluted. Essentially, it constructs upwards from processed data (e.g., reconstructs an image from a label).
C. **Generative Adversarial Network** - A large family of networks composed of two separate networks, a generator and a discriminator. The generator is trained to generate realistic data, while the discriminator is trained to differentiate between generated data from actual samples. Essentially, in this approach, the objective is for the generator to generate such realistic data that the discriminator cannot tell it apart from samples.
D. **Autoencoders** - A family of networks is trained in an unsupervised manner. The autoencoder aims to learn how to robustly represent the original dataset, oftentimes in smaller dimensions, even in the presence of noise. Autoencoders are composed of multiple layers, and it can be divided into two main parts: the encoder and the decoder. The encoder maps the input into the representation and the decoder uses the representation to reconstruct the original input.
E. **Deep Belief Network** - A family of generative networks that are composed of multiple layers of hidden units, in which there can be connections between layers but not within units within layers. Deep belief networks can be conceived as being composed of multiple simpler networks, where each subnetwork's hidden layer acts as a visible layer to another subnetwork.

Learning Classes

A. **Supervised Learning** - Training data is provided when fitting the model. The training dataset is composed of inputs and expected outputs. Models are tested by making predictions based on inputs and comparing them with expected outputs.
B. **Unsupervised Learning** - No training data is provided to the model. Unsupervised learning relies exclusively on inputs. Models trained using unsupervised learning are used to describe or extract relationships in image data, such as clustering or dimensionality reduction.
C. **Reinforcement Learning** - The learning process occurs in a supervised manner, but not through the use of static training datasets. Rather, in reinforcement learning, the model is directed towards a goal, with a limited set of actions it may perform, and model improvement is obtained through feedback. The learning itself occurs exclusively through feedback obtained based on past action. This feedback can be quite noisy and delayed.
D. **Hybrid Learning Problems**
   Semi-Supervised Learning - Semi supervised learning relies on training datasets where only a small percentage of the training dataset is labeled, with the remaining images having no label. It is a hybrid in between supervised and unsupervised learning, since the model has to make effective use of unlabeled data while relying only partially on labeled ones.
   Self-Supervised Learning - Self supervised learning uses a combination of unsupervised and supervised learning. In this approach, supervised learning is used

to solve a pretext task for which training data is available (or can be artificially provided), and whose representation can be used to solve an unsupervised learning problem. Generative adversarial networks rely on this technique to learn how to artificially generate image data.

Other learning Techniques

A. **Active Learning** - During active learning, the model can query the user during the learning process to require labels for new data points. It requires human interaction, and it aims to being more efficient about what training data is used by the model
B. **Online Learning** - Online learning techniques are often used in situations where observations are streamed through time and in which the probability distribution of the data might drift over time. In this technique, the model is updated as more data becomes available, allowing the model itself to change through time.
C. **Transfer Learning** - Transfer learning is a useful technique when training a model for a task that is related to another task for which a robust model is already available. Essentially, it treats the already robust model as a starting point from which to train a new model. It greatly diminishes the training data needs of supervised models and it is, therefore, used when the available training data is limited.
D. **Ensemble Learning** - As mentioned in the main text, ensemble learning refers to a learning technique in which multiple models are trained either in parallel or sequentially and the final prediction is the result of the combination of the predictions generated by each component.