ChemGrapher: Optical Graph Recognition of Chemical Compounds by Deep Learning

Martijn Oldenhof
ESAT - STADIUS
KU Leuven
Leuven, 3001, Belgium
martijn.oldenhof@esat.kuleuven.be

Adam Arany
ESAT - STADIUS
KU Leuven
Leuven, 3001, Belgium
adam.arany@esat.kuleuven.be

Yves Moreau
ESAT - STADIUS
KU Leuven
Leuven, 3001, Belgium
moreau@esat.kuleuven.be

Jaak Simm
ESAT - STADIUS
KU Leuven
Leuven, 3001, Belgium
jaak.simm@esat.kuleuven.be

Abstract

In drug discovery, knowledge of the graph structure of chemical compounds is essential. Many thousands of scientific articles in chemistry and pharmaceutical sciences have investigated chemical compounds, but in cases the details of the structure of these chemical compounds is published only as an images. A tool to analyze these images automatically and convert them into a chemical graph structure would be useful for many applications, such drug discovery. A few such tools are available and they are mostly derived from optical character recognition. However, our evaluation of the performance of those tools reveals that they make often mistakes in detecting the correct bond multiplicity and stereochemical information. In addition, errors sometimes even lead to missing atoms in the resulting graph. In our work, we address these issues by developing a compound recognition method based on machine learning. More specifically, we develop a deep neural network model for optical compound recognition. The deep learning solution presented here consists of a segmentation model, followed by three classification models that predict atom locations, bonds and charges. Furthermore, this model not only predicts the graph structure of the molecule but also produces all information necessary to relate each component of the resulting graph to the source image. This solution is scalable and could rapidly process thousands of images. Finally, we compare empirically the proposed method to the well-established tool OSRA [8] and observe significant error reductions.

1 Introduction

Knowledge of the chemical structure of compounds is central in drug discovery because this structure determines the properties of the compound. It is for example used for drug candidate selection. Because billions of euros of research and development investment are needed to successfully bring a new drug to the market, any tool that improves the drug candidate selection process would have a significant pharmaceutical impact.

Although chemical structures, which are the familiar graph drawings of molecules, do lose some information about the electronic structure of a molecule (which is actually responsible

Preprint. Under review.

for its chemical properties), they are powerful and effective abstractions. To query such structures or apply machine learning, we need to start from a well-structured data set encoding the graph representation of the chemical structure. This encoding step, which is usually less flexible than an arbitrary drawing, might lose also some information about the chemical structure, but it will provide a solid starting point for further automated processing. Interesting formats for representing chemical compounds are for example SMILES [27] or MOLfile [6], which contains all necessary information to build the complete molecule graph structure. Using these formats, it would for example be possible to query documents for specific patterns in chemical compounds. However, such encodings still remains somewhat cumbersome and are not yet systematically available, in particular for unstructured legacy data.

Thousands of scientific publications describe new chemical compounds and investigate their properties. However, the structure of these chemical compounds are usually described in the publication only as an image. This means that today a rich source of data, which would be extremely valuable to develop novel machine learning approaches or simply query documents more accurately, is largely under-exploited. It is therefore important to convert images of chemical structures into these formats. A few tools for recognizing graph structures from chemical compound images are available, such as OSRA [8] and others like ChemReader [18], Kekule [17], CLiDE Pro [25], and the work of M. Sadawi et al.. However, we observe that, using these tools, bond multiplicity and stereochemical information are sometimes lost. Those tools are mainly expert systems using different techniques, such as image processing, optical character recognition, hand-coded rules, or sophisticated algorithms. Modifying or further improving these tools requires a lot of effort. A tool based on machine learning, which learns directly from training data, would be most valuable. Such a tool could potentially become more accurate than existing methods and its performance could be improved by increasing the size and the diversity of the data sets, instead of having to modify its code.

Therefore, we propose a new *data-driven* machine learning based tool that can learn from only image data to recognize the chemical structure graph given an image of a chemical structure. The core of the tool is a deep learning model. In the work of Staker et al., another deep learning model was also proposed. However, there the output is only a text-sequence representing the graph. In our approach, we focus on directly predicting the graph structure (*i.e.*, identifying all the nodes and the edges and their labels). The positions of these nodes and edges in the resulting graph would correspond with the positions in the original image of the chemical structure. The resulting graph can be later translated to any format (*e.g.*, SMILES).

In the next sections, we will describe the method and the neural networks used, and also how the different networks interact. Next, we describe the data sets used for training. Then we focus on the performance and scalability of our method, and conclude with possible future work.

2 Related work and Background

A deep convolutional neural network [15] is the type of network most often used for image recognition. These convolutional neural networks can be split into two main types: (1) image segmentation networks and (2) classification networks. We combine both approaches in our graph recognition tool.

2.1 Image segmentation

Our work builds upon the recent developments in image segmentation. Different machine learning approaches can be used for the segmentation of images. One well-established approach is U-Net [20]. This approach uses a network that combines a contracting path and an expanding path. Several other works were based on the U-net approach, such as Jansson et al., where a U-net is used to extract the vocal component from music. Other works expanded this U-net approach, such Çiçek et al., which generalizes the U-net approach to 3D images. Another approach is to make use of dilated convolutions [30] all stacked without loss of resolution. An advantage of dilated convolutions is that the receptive field can grow

exponentially by increasing the dilation in the dilated convolutional operator. This would be computationally more interesting than using multiple convolutions or larger kernels.

2.2 Image classification

There has been a trend to create deeper and deeper neural networks to improve performance for the classification of images. However, deeper networks are more difficult to train. To improve the training capabilities of such deep networks, methods such as residual neural networks (Resnet) [11] were developed. It has also been shown that these residual neural network are comparable in performance and behavior with an ensemble of more shallow networks [26]. It is also worth mentioning the work of Zhang et al. where the concept of Resnet is combined with the concept of U-net for the localization of roads on aerial images.

2.3 Drug discovery and machine learning

There are several stages in the process of drug discovery. The stages go from basic research and drug candidate selection to the development phase, clinical trials, and finally production. As development progresses further and sunken costs increase, the cost of failure of a project thus increases. "Fail early" is thus important to contain the costs of drug discovery. Predicting risks of failure later in the discovery process (for example, by predicting toxicity for a compound) without draining the pipeline (enough candidate compounds need to remain available) is essential. Machine learning techniques can be used in all stages of drug discovery. Chen et al. gives a good overview of the recent use of deep learning in drug discovery. We would like to highlight some of these recent applications, which we find interesting in the context of our graph recognition tool.

In the first place, there is the work from Xu et al. and Gómez-Bombarelli et al., where an unsupervised method is used to extract features from SMILES input data. SMILES(Simplified Molecular Identification and Line Entry System) [27] is a text representation of a chemical compound following specific syntactic rules. The unsupervised learning methods in both works are based on the auto-encoder principle. This feature vector can then be used as input to a supervised method to learn to predict molecular properties.

Another interesting method to predict molecular properties of a chemical compound is to use the neural graph fingerprint presented in Duvenaud et al.. The neural graph fingerprint is a way to represent and encode a chemical compound. Here, a convolutional neural network takes these graphs as input and is trained to predict molecular properties. Similarly, in Kearnes et al., Coley et al., Simm et al. and Pires et al., a machine learning model takes a molecular graph as input.

Large amount of data is needed to use or train the models mentioned above. It is not always easy to find this data. This is where our tool is useful, by extracting graph representations of chemical compounds directly from images. It is also worth mentioning the work presented in B. Goh et al., where no representation of the chemical compound is needed. In this work, a machine learning model is trained to predict molecular properties directly from images from chemical structures.

2.4 Stereoisomerism

Stereochemical information can also be encoded in a 2D representation of a molecule. This stereochemical information is important to differentiate molecules with the same molecular formula but with a different spatial orientation. To encode this central chirality, different type of lines are used to represent bonds in the 2D representation of a molecule: solid lines, wedge-shaped lines or dashed lines [23]. It is important that this information is captured by our graph recognition tool.

3 Problem Statement

In this section, we formulate our learning task. The goal of the proposed method is to learn a function that maps an image $\mathbf{x}$ to its graph representation G.

Definition 1 $\mathbf{x} \in \mathbb{R}^{U \times V}$ *represents a single-channel 2D image with dimensions* $U \times V$.

Generalizations to multiple channels is straightforward, if colored inputs are available.

Definition 2 $G = (V, E)$ *represents a graph with vertices* V *and edges* E.

For our graph recognition tool to work, we need to learn the following function:

$$g : \mathcal{X} \to \mathcal{G}. \tag{1}$$

This function will map an 2D input image of a chemical structure to the graph representation of the molecule. To learn this function, we make the following assumption for the training data set:

Assumption 1 *We assume we have the knowledge about the location of every node in the graph represented by* $\mathbf{L} \in \mathbb{N}^{i \times 2}$*, with* i *number of nodes in the graph for every image in our labeled training data set. We also assume the knowledge of all inter-node connections in the graph (edges) for all labeled images represented by* $\mathbf{C} \in \{0, 1\}^{i \times i}$.

Definition 3 *The training data set* $\mathcal{D}$ *is then defined as* $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{L}_i, \mathbf{C}_i)\}_{i=1}^{N_{\text{samples}}}$.

4 Model

To solve the problem statement defined in the previous section, we build a machine learning model. The model is split up in different learning tasks, which will be defined here.

4.1 First task: segment type segmentation

The first learning task is to learn to segment a 2D image of a chemical structure in different segments, where each segment can represent the location of a specific atom, charge or bond type in the image. The image was already defined in previous section. However, the segmentation of this image will be defined here.

Definition 4 $\mathbf{S}^a \in \mathbb{R}^{U \times V \times n_a}$, $\mathbf{S}^b \in \mathbb{R}^{U \times V \times n_b}$, $\mathbf{S}^c \in \mathbb{R}^{U \times V \times n_c}$ *represent the atom type, bond type and charge segmentation of an image.* U *and* V *are the same as in the input image while* n_a*,* n_b *and* n_c *respectively are the number of atom types, bond types and charges (including the empty atom, charge and bond type) present in the compound.*

To perform image segmentation, we need to learn the following function

$$s : \mathcal{X} \to \mathcal{S}^a, S^b, S^c. \tag{2}$$

To learn this function, we need to label the training elements.

Definition 5 *Let* $\mathbb{I}_n = \{1, ..., n\} \subseteq \mathbb{N}$ *then we define* $\mathbf{L}^a \in \mathbb{I}_{n_a}^{U \times V}$, $\mathbf{L}^b \in \mathbb{I}_{n_b}^{U \times V}$, $\mathbf{L}^c \in \mathbb{I}_{n_c}^{U \times V}$ *which represent the pixelwise true labels.* U *and* V *are the same as the input image.* n_a*,* n_b *and* n_c *are respectively the number of atom types, bond types and charge (including the empty atom, charge and bond type). The value of every element represent the atom, charge and bond type to which the corresponding pixel belongs.*

Once the true labels for the training data have been defined, we can define the loss function for training. Here, we will use the Cross Entropy Loss, which is defined as

$$H(y, \hat{y}) = -\sum_{i}^{N} y_i \log \hat{y}_i, \tag{3}$$

where y_i is the true probability distribution of the true labels, $\hat{y}_i$ is the estimated probability distribution of the labels, and N is the number of different classes. In the case of atom type segmentation $\mathbf{S}^a$, the cross entropy loss is calculated and summed for every pixel prediction (so fixing u and v) in the following way, taking into account that $\mathbf{S}^a$ is not a probability distribution:

$$\text{Loss}_a = -\sum_{u=1}^{U} \sum_{v=1}^{V} \log \left(\frac{\exp(\mathbf{S}^a_{u,v,\mathbf{L}^a_{u,v}})}{\sum_j \exp(\mathbf{S}^a_{u,v,j})} \right). \tag{4}$$

The loss (Loss_b, Loss_c) in the case of bond type segmentation $\mathbf{S}^b$ and charge segmentation $\mathbf{S}^c$ is calculated similarly. The total loss is the sum of all partial losses:

$$\text{Loss}_{total} = \text{Loss}_a + \text{Loss}_b + \text{Loss}_c. \tag{5}$$

4.2 Second task: segment type classification

A second learning task is necessary to build a final graph. This learning task classifies parts of the segmented image into the different possible atom, bond and charge types. One part of the input used in this learning task is defined in the following way:

Definition 6 $\mathbf{S}^{a_{cut}} \in \mathbb{R}^{K \times L \times n_a}$, $\mathbf{S}^{b_{cut}} \in \mathbb{R}^{K \times L \times n_b}$, $\mathbf{S}^{c_{cut}} \in \mathbb{R}^{K \times L \times n_c}$ *where* $K, L = 2$ *times the regular bond length in a 2D image of a chemical structure and* n_a, n_b *and* n_c *are respectively the number of different atom types, bond types or charges (including the empty types). The tensors* $\mathbf{S}^{a_{cut}}$, $\mathbf{S}^{b_{cut}}$ *and* $\mathbf{S}^{c_{cut}}$ *represent a cut-out of the tensors* $\mathbf{S}^a$, $\mathbf{S}^b$ *or* $\mathbf{S}^c$.

Another part of the input used in this learning task is defined as:

Definition 7 $\mathbf{x}^{a_{cut}} \in \mathbb{R}^{K \times L}$, $\mathbf{x}^{b_{cut}} \in \mathbb{R}^{K \times L}$, $\mathbf{x}^{c_{cut}} \in \mathbb{R}^{K \times L}$ *represent the cut-outs of the original 2D image* $\mathbf{x} \in \mathbb{R}^{U \times V}$ *where* $K, L = 2$ *times the regular bond length in a 2D image of a chemical structure.*

Next, we define the output used in this learning task.

Definition 8 $A \in \mathbb{R}^{n_a}$, $B \in \mathbb{R}^{n_b}$, $C \in \mathbb{R}^{n_c}$ *are vectors of dimension* n_a, n_b *and* n_c *where* n_a, n_b *and* n_c *are respectively the number of different atom types, bond types or charges (including the empty types). The vectors* A, B *and* C *represent respectively the atom type, bond type and charge predictions.*

With these definitions, we can now also define the functions to be learned in this task:

$$c_A : \mathcal{S}^{a_{cut}}, \mathcal{X}^{a_{cut}} \rightarrow \mathcal{A}, \tag{6}$$

$$c_B : \mathcal{S}^{b_{cut}}, \mathcal{X}^{b_{cut}} \rightarrow \mathcal{B}, \tag{7}$$

$$c_C : \mathcal{S}^{c_{cut}}, \mathcal{X}^{c_{cut}} \rightarrow \mathcal{C}. \tag{8}$$

To learn these functions, we need the labels of the training data.

Definition 9 *Let* $\mathbb{I}_n = \{1, ..., n\} \subseteq \mathbb{N}$ *then we define* $L_A \in \mathbb{I}_{n_a}$, $L_B \in \mathbb{I}_{n_b}$ *and* $L_C \in \mathbb{I}_{n_c}$ *respectively represent the label of atom type, bond type and charge of each training element.* n_a, n_b *and* n_c *are respectively the number of atom types, bond types and charge (including the empty atom, charge and bond type).*

Finally, we also define the loss function used in the training phase for learning funtion c_A.

$$\text{Loss}_{c_A} = -\log\left(\frac{\exp(A_{L_A})}{\sum_j \exp(A_j)}\right) \tag{9}$$

The loss functions (Loss_{c_B}, Loss_{c_C}) for the functions c_B and c_C can be defined similarly.

5 Graph Building Algorithm

Once we have learned the functions described in the previous section we need an algorithm to combine the outputs of these functions and build up a final graph structure. We propose the algorithm defined in Algorithm 1:

Algorithm 1: Graph building algorithm

Data: Image tensor $\mathbf{x}$
Result: Graph G
$\mathbf{S}^a$,$\mathbf{S}^b$,$\mathbf{S}^c = s(x)$
atomCandidates = generateAtomCandidates($\mathbf{S}^a$)
$V = []$
for atomCand ***in*** atomCandidates **do**
  $\mathbf{S}^{a_{cut}}$,$\mathbf{x}^{a_{cut}}$,$\mathbf{h}_L^a$ = cutAtomCand(atomCand, $\mathbf{S}^a$, $\mathbf{x}$)
  $\mathbf{S}^{c_{cut}}$,$\mathbf{x}^{c_{cut}}$,$\mathbf{h}_L^c$ = cutAtomCand(atomCand, $\mathbf{S}^c$, $\mathbf{x}$)
  $A = c_A(\mathbf{S}^{a_{cut}}, \mathbf{x}^{a_{cut}}, \mathbf{h}_L^a)$
  $C = c_C(\mathbf{S}^{c_{cut}}, \mathbf{x}^{c_{cut}}, \mathbf{h}_L^c)$
  if isNotEmptyAtomType(A) **then**
    V.appendAtom(A, C, atomCand)
  end
end
bondCandidates = generateBondCandidates(V)
$E = []$
for bondCand ***in*** bondCanidates **do**
  $\mathbf{S}^{b_{cut}}, \mathbf{x}^{b_{cut}}, \mathbf{h}_L^b$ = cutBondCand(bondCand, $\mathbf{S}^b$, $\mathbf{x}$)
  $B = c_B(\mathbf{S}^{b_{cut}}, \mathbf{x}^{b_{cut}}, \mathbf{h}_L^b)$
  if isNotEmptyBondType(B) **then**
    E.appendBond(B, bondCand)
  end
end

The proposed algorithm 1 will first apply the segmentation function s to the input image. Next, given the segmentation $\mathbf{S}^a$, candidate locations atomCandidates will be generated by generateAtomCandidates. Given these candidate locations, the nodes V of the graph can be build in an iterative way. For this purpose, the segmentations $\mathbf{S}^a$ and $\mathbf{S}^c$ can be cut (cutAtomCand) into smaller segments $\mathbf{S}^{a_{cut}}$ and $\mathbf{S}^{c_{cut}}$ thanks to every candidate location atomCand. At the same time the original image $\mathbf{x}$ is also cut (cutAtomCand) into smaller parts $\mathbf{x}^{a_{cut}}$ and $\mathbf{x}^{c_{cut}}$. Extra highlights are also created $\mathbf{h}_L^a$ and $\mathbf{h}_L^c$ which highlight the candidate location to be classified. Then, the classification functions c_A and c_C are applied to determine what kind of atom and charge is located at the candidate location. If the candidate location is not empty (isNotEmptyAtomType), the location, type A and charge C will be added to the list of nodes V. Later, the algorithm will use these nodes V to build the edges E of the graph G. For this, it first will need to generate (generateBondCandidates) the candidate bond locations (bondCandidates). Similarly, as for the nodes, the bonds E of the graph can be built in an iterative way. For this purpose, the segmentation $\mathbf{S}^b$ can be cut (cutBondCand) into a smaller segment $\mathbf{S}^{b_{cut}}$ thanks to every candidate bond location bondCand. At the same time the original image $\mathbf{x}$ is also cut (cutBondCand) into a smaller part $\mathbf{x}^{b_{cut}}$. Extra highlights are also created $\mathbf{h}_L^b$ which highlight the bond location to be classified. Finally, the classification function c_B is applied to determine the type of bond located at the candidate bond location. If the candidate bond location is not empty (isNotEmptyBondType), the location and type B will be added to the list of bonds E.

6 Deep learning implementation

The method we use in the graph recognition tool is a combination of different convolutional neural networks [15]. First, we have a semantic segmentation network using the Dense Prediction Convolutional Network [21, 30] followed by three classification networks. The output of the segmentation network is part of the input of the other classification networks.

Table 1: Summary of the layers of the segmentation network

Layer	Kernel	Nonlinearity	Padding	Dilation
conv1	3x3	ReLU	1	no dilation
conv2	3x3	ReLU	2	2
conv3	3x3	ReLU	4	4
conv4	3x3	ReLU	8	8
conv5	3x3	ReLU	8	8
conv6	3x3	ReLU	4	4
conv7	3x3	ReLU	2	2
conv8	3x3	ReLU	1	no dilation
last	1x1	none	no padding	no dilation

6.1 Semantic segmentation network s

Before feeding the image to the segmentation network s, the image is preprocessed to a binary black and white image. The output of the segmentation network are different channels predicting for every pixels in the image the class the pixels belong to. The possible classes represent the different atom types, bond types and charges. For the implementation of this network, we build on the concept of dilated convolution described in Yu and Koltun.

6.1.1 Network architecture

The network has 8 3x3 convolutional layers from which 6 layers make use of dilation. All convolutional layers are followed by a Rectified Linear Unit (ReLU). The last layer is a linear layer. Padding is used so that the resolution of the channels does not change. The padding and dilation for the different convolutional layers are summarized in Table 1.

6.2 Classification networks

For the atom location, the bond prediction and the charge prediction we use three separate classification networks. All three networks use part of the output of the segmentation networks in their input.

6.2.1 Atom location prediction c_A

For the atom location prediction, part of the output the segmentation network ($\mathbf{S}^a$) is used. This output contains the segmentation of only the atoms in the image. Next, the original binary image (x) is also used for the input. Finally, candidate locations are created (generateAtomCandidates) to spot the part of the image we want to classify. This is also formatted as input and fed into the network. Depending on the candidate location, the inputs can be reduced (cutAtomCand) so that only the immediate region of the candidate location is included ($\mathbf{S}^{a_{cut}}$). This will reduce the computational cost and speed up the learning of the network. This process is illustrated in Figure 1.

The output of this network is a vector (A) with every element representing the prediction of the network. The size of this vector is the number of different atom classes, plus one to represent the empty class (no atom). For every image segmentation, this network has to run several times to classify all candidate locations to get all atom predictions in the original image.

6.2.2 Bond prediction c_B

For the bond prediction network, we apply a similar strategy as illustrated in Figure 2. This time another part of the output the segmentation network is used ($\mathbf{S}^b$). This output contains the segmentation of only the bonds in the image. Every type of bond is represented in the segmentation as a rectangle. For stereo bonds, we use two rectangles to encode the direction of the bond. Next, as in the atom prediction network, the original binary image (x) is used also for the input of the bond prediction network. Finally, for the bond prediction, we also

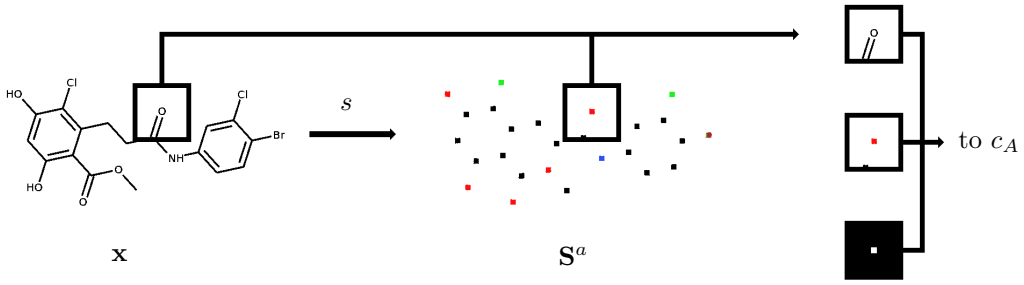

Figure 1: To build the input for the atom classification network (c_A), the output of the segmentation network $\mathbf{S}^a$ is cut (cutAtomCand) to feed it in to the atom classification network. This cut-out is shown in the middle ($\mathbf{S}^{a_{cut}}$). To this, we also add part of the original image ($\mathbf{x}^{a_{cut}}$) together with highlighting the candidate location ($\mathbf{h}_L^a$) of the atom type to be classified. The complete input for c_A is shown on the right.

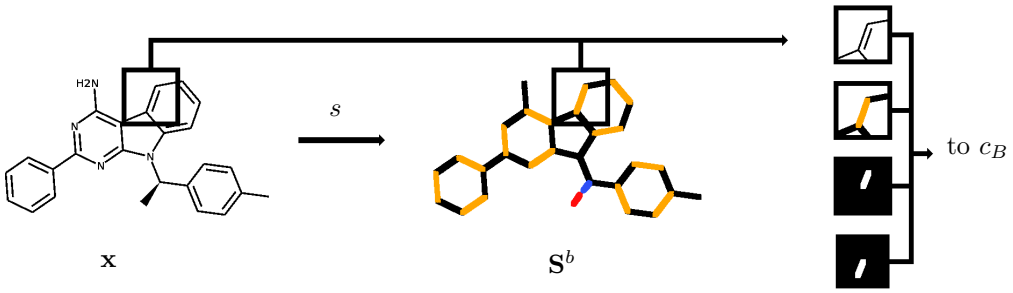

Figure 2: To build the input for the bond classification network (c_B), part of the segmented image ($\mathbf{S}^b$) is cut out (cutBondCand). This cut-out ($\mathbf{S}^{b_{cut}}$) is shown in the middle of the figure. To this, we also add part of the original binary input image ($\mathbf{x}^{b_{cut}}$) and the candidate bond location ($\mathbf{h}_L^b$) encoded in two parts. The complete input for c_B is shown on the right.

need to encode candidate locations for the bond predictions. This time, as opposed to the atom prediction network, we will use two parts. One rectangle represents the first part of the bond connected with the first atom and another rectangle represents the second part connected with the second atom. The rectangles meet in the middle. By using two rectangles we can encode the direction of the bond which is necessary to predict the stereoisomeric bond direction. These candidate locations are generated (generateBondCandidates) from the predictions from the atom location network. Moreover, depending on these locations, we can cut out (cutBondCand) the inputs again so that only the immediate region of these candidate pairs is fed into the network.

Like with the atom prediction network, the output is again a vector (B) with every element representing the prediction of the network. This time the vector size is the number of different bond classes, plus one to represent the empty class (no bond).

6.2.3 Charge prediction c_C

As with the bond prediction and the atom prediction network a similar strategy is used illustrated in Figure 3. Again a part of the output the segmentation network is used ($\mathbf{S}^c$). This output contains the segmentation of only the charges in the image. Every charge is represented by a rectangle located on the location of the atom it applies to. Depending on the candidate location generated by generateAtomCandidates, the inputs can be again reduced (cutAtomCand) so that only the immediate region of the candidate location is included ($\mathbf{S}^{c_{cut}}$). Again the original image (x) is fed together with the candidate location as input to the charge prediction network.

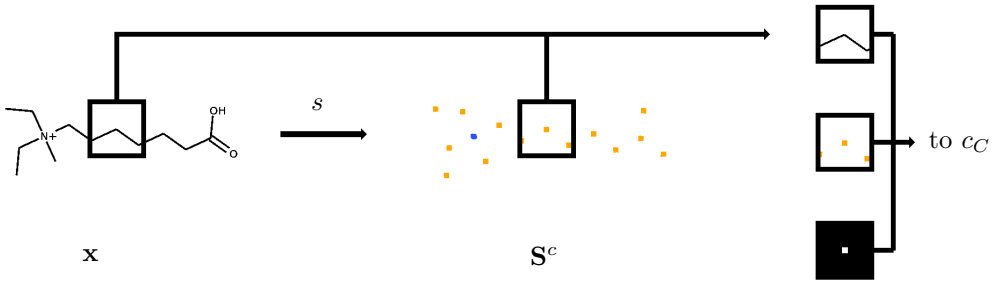

Figure 3: To build the input for the charge classification network (c_C), the output of the segmentation network $\mathbf{S}^c$ is cut (cutAtomCand) to feed it in to the charge classification network. This cut-out is shown in the middle ($\mathbf{S}^{c_{cut}}$). To this, we also add part of the original image ($\mathbf{x}^{c_{cut}}$) together with highlighting the candidate location ($\mathbf{h}_L^c$) of the charge type to be classified. The complete input for c_C is shown on the right.

Table 2: Different layers in the classification network

Layer	Kernel	Nonlinearity	Padding	Dilation
depthconv1	3x3	ReLU	1	no dilation
conv2	3x3	ReLU	2	2
conv3	3x3	ReLU	4	4
conv4	3x3	ReLU	8	8
conv5	3x3	ReLU	1	no dilation
maxpool	124x124	None	no padding	no dilation
last	1x1	None	no padding	no dilation

Like with the atom and bond prediction network, the output is again a vector (C) with every element representing the prediction of the network. This time the vector size is the number of different charge classes, plus one to represent the empty class (no charge).

Once we have the bond predictions together with the atom and charge predictions, we can build the graph structure of the complete molecule.

6.2.4 Network architecture

The three classification networks have similar layer structures. There are 5 convolutional layers where 3 of them are dilated and 1 of them (the first one) is actually a depthwise separable convolution [3]. After the convolutional layers, there is always a ReLU layer. The last layer is a linear layer and the layer before that is a max pool layer. All layers are summarized in Table 2.

7 Data Sets

To build our data sets for the segmentation network and the classification networks we download and split different chemical structures in SMILES format from the ChEMBL [9] database. The around 1.9 million chemical structures are splitted in 4 parts:

- a training pool for the segmentation network of 1.5 million chemical structures,
- a pool of 300K chemical structures used for the validation of the segmentation network and training of classification networks,
- a pool of around 35K chemical structures for the validation of the classification networks and
- another pool of 35K chemical structures for testing the overall performance.

From these pools we can sample the actual data sets for our different networks. By sampling we can control the relative frequency of different atom types and bond types in the actual data sets. This is important for the performance of our networks, due to data imbalance, on the different atom types and bond types as we will see in the next section.

7.1 Segmentation Data Set

For the training of the segmentation network (s), we need 2D images of chemical structures together with pixelwise labeled target values ($\mathbf{L}^a$, $\mathbf{L}^b$ and $\mathbf{L}^c$). This type of data set is not available as far as we know, so we need to construct this data set ourselves. Labeling thousands of 2D images of chemical structures pixelwise by hand is moreover not feasible. We thus construct an automatic procedure to generate this data set. For the training data set we sample around 114K chemical compounds in SMILES format from the ChEMBL training pool in a way that every atom type is present in at least 1000 chemical compounds. Using RDkit [14] in Python, we create the images starting from the SMILES. Furthermore, to create the labelling, we make some modifications in the code of RDkit at the drawing time of the image, so that it additionally produces the necessary labeling information needed to create our data set. For the validation dataset the same procedure can be used.

7.2 Atom prediction data set

Once the segmentation network is trained and validated we can sample from the ChEMBL data a new data set to feed into the segmentation network. The output of these runs are saved to create the input data set for the next classification networks. As already explained in the previous section, the atom prediction network (c_A) additionally expects as input the candidate locations to classify. For the training and validation data sets of the classification network c_A, we generate candidate locations based on the true atom location values, but also add locations where no atom is located for the prediction of the empty class. For these locations, we take the middle point of every bond in the data set. As we know that no atom is located in the middle of a bond, these locations can be used for the empty values in the data sets.

7.3 Bond prediction data set

For the bond prediction network (c_B), we apply a similar technique. In addition to the inputs from previous segmentation network, the bond prediction network expects the candidate bond locations. For the training and validation data of c_B, we generate these candidate locations by going over all possible combinations of pairs of atoms in a molecule within the range of less than two times the bond length. Sometimes there is a bond between a generated pair of atoms and then the data set item is labeled with the type of bond. If there is no bond between a pair of atoms, the item is labeled as empty.

7.4 Charge prediction data set

For the charge classification network the same data sets as the atom prediction data sets can be used except for the labels. Instead of the atom types the labels would now be the charge (including empty charge) of the atom candidate.

8 Experiments and results

For validation we sample new data sets from the validation ChEMBL pools for the different networks. For the segmentation network we sample around 12K chemical structures. For the validation of the classification networks less chemical structures are needed, so we only sample around 450 chemical structures. Starting from these 450 chemical structures we generate atom candidates and bond candidates. This results in a data set of around 27K of atom candidates for atom type and charge classification networks and a data set of around 55K of bond candidate locations for the bond classification network. With these validation data sets we measure the performance on the different networks.

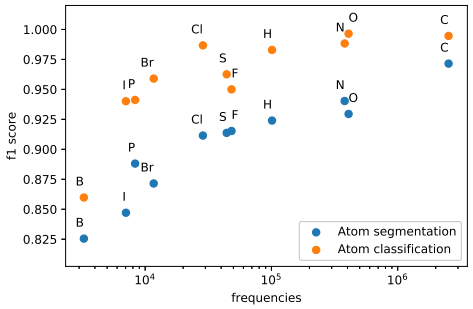

(a) Atom prediction performance (s_A and c_A)

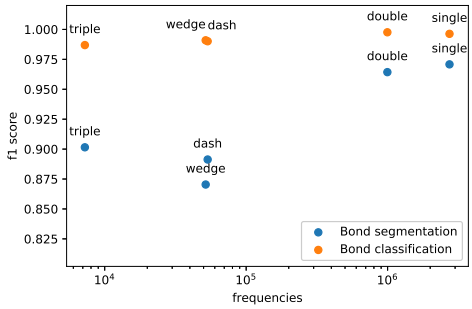

(b) Bond prediction performance (s_B and c_B)

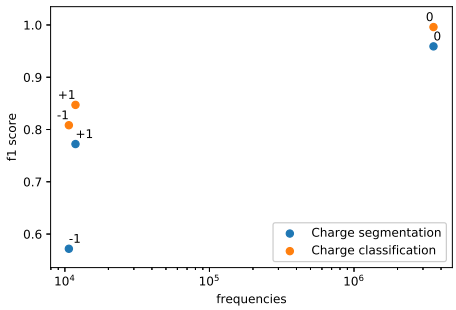

(c) Charge prediction performance (s_C and c_C)

Figure 4: F1 score for segmentation and classification networks. There is clearly a correlation between the performance of the networks on the different prediction types and the frequency of the specific type in the training data set. The classification networks perform significantly better than the segmentation networks.

8.1 Performance of segmentation network

For the segmentation network (s) we measure the F1 score [28] for all the pixel predictions for the different atom, bond and charge types. The F1 score takes into account both precision and recall equally. If we compare the F1 score with the frequency of the different atom, bond and charge types in the training data set we clearly see a correlation. The results are summarized in Figure 4 .

8.2 Performance of classification networks

For the classification networks we again use the F1 score to measure the performance for the atom, bond and charge type classifications. Again we see a correlation between the F1 score and the frequency of the different types in the training data set. We can also empirically see that the F1 score for the classification networks is significantly higher than for the segmentation networks. So the classification networks can do a good job even when the segmentation is not perfect. The performance of these classification networks have to be very good as for every graph prediction tens of bond and atom classifications have to made and this would otherwise degrade the overall accuracy rapidly. The results are summarized in Figure 4 .

8.3 Overall graph accuracy

Now that we know the performance of the different parts, we can combine those building blocks and measure the overall accuracy of the resulting graph predictions. As already mentioned in a previous section, the segmentation network and classification networks should be used as presented in algorithm 1 in order to build the resulting graph. Images in 3 different

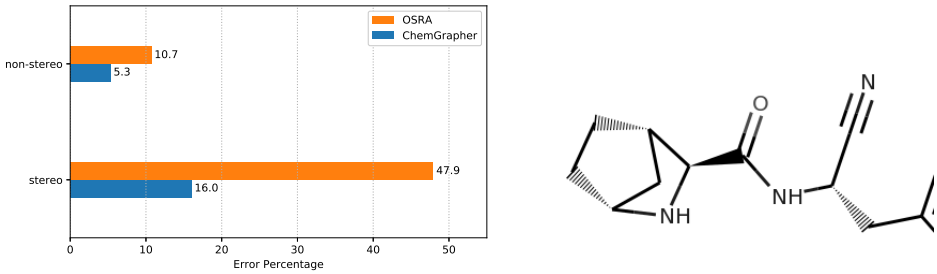

(a) Error rate for style 1

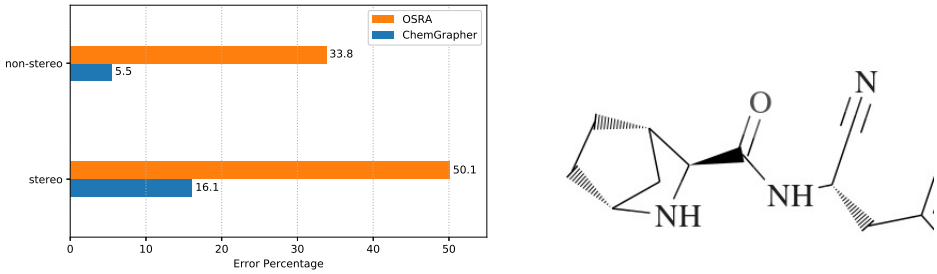

(b) Error rate for style 2

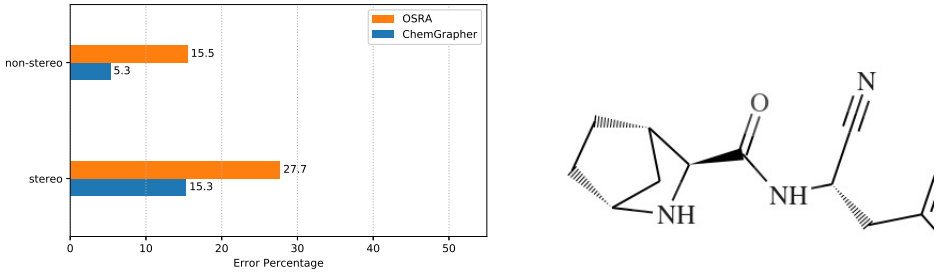

(c) Error rate for style 3

Figure 5: The graph accuracy of our tool compared with OSRA measured on images generated in different styles. For each style 2 experiments are performed: once with images without stereo-chemical information and once with images with stereo chemical information. On the left we observe for each style the results on the error rate and on the right we observe for each style an example image in that specific style. For all styles we measure a lower error rate for our tool ChemGrapher compared to OSRA.

styles are generated and for every style we generate 2 sets where 1 set only has images without stereo chemical information encoded in the compounds while the other set has images where all compounds have stereo chemical information encoded. This results in 6 sets of each 1000 images to measure the performance on our tool ChemGrapher. If we count at least one mistake in the resulting graph we count the graph prediction as incorrect. The same sets we also use to measure the performance on OSRA to compare. The results are summarized in Figure 5 . On all sets we observe a a higher accuracy on our tool ChemGrapher compared to OSRA.

8.4 Case Study: Performance on Journal Article Images

The idea of ChemGrapher is to use it on images in journal articles therefore we would also like to know how well this tool performs on such images. As this kind of data set is not available we decided to build one manually. So we cut out images from journal articles about

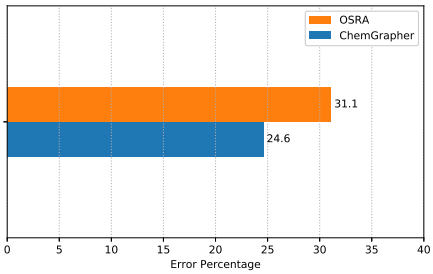

Figure 6: Error rate of our tool ChemGrapher on test set of journal article images compared with OSRA. From the errors we learn there is still room for improvement in future work.

chemical compounds, preprocess them to the correct input format and feed them to our tool. We evaluate the resulting graph manually on correctness to measure the accuracy. If we count at least one mistake in the resulting graph we categorize the prediction as incorrect. The same procedure was executed for OSRA for comparison. The results of this experiment are summarized in Figure 6. Thus, out of total of 61 images we tried on ChemGrapher, 46 were correctly predicted while OSRA predicted 42 images correctly. However, we can also observe that ChemGrapher clearly has better performance on images of compounds with only carbon atoms compared to OSRA. For these compounds typically no letters appear in the image. Another observation we make is that ChemGrapher still has some issues related when thick lines are used to depict the bonds. We set this as a target for our future work.

9 Future Work

To train the segmentation network, we need a pixelwise labeled data set. However, this kind of data set is not always available. We thus created this data set with RDkit. However the consequence is that the format of the input image is somewhat biased. We have seen in the case study that ChemGrapher performs reasonably although not equally well on real images. To handle other kind of image formats, it might be difficult to find a pixelwise labeled data set to retrain our networks. Therefore, future work could focus on building a method that can learn from data that is not labeled pixelwise. The data would only offer a way to verify if the resulting graph is correct or not. We could consider this as an instant of weakly supervised learning.

10 Conclusion

We presented a method to recognize the graph structure of molecules from 2D images of chemical structures using deep learning. This method learns a model directly from data. We have seen that careful data preparation is crucial. Care should be taken to have a balanced data set for the different classes of atoms and bonds. However, even with an imperfectly balanced data set, our deep learning methods give very good results. One thing that is important for our method to work is the fact that the classification networks need to have an almost perfect accuracy. While the segmentation network can tolerate some errors, for the classification networks every drop in accuracy can have dramatic results on the overall accuracy. The performance is also clearly better than the well known tool OSRA [8] and also provides us more detailed information about the resulting graph. For our deep learning method to learn accurately, we also had to implement an automatic procedure to pixelwise label 2D images of chemical structures. We described how we modified the code of RDkit for this purpose. In fact, this pixelwise labelling of images for the segmentation is actually key to linking the atoms and bonds in the resulting graph back to the source image. This makes this deep learning method non-black box or interpretable. In the context of drug discovery, such tools are important. In fact, in general we see that machine learning is gaining importance in this area and that it contributes to improving the quality of the drug discovery process.

Acknowledgments

MO, AA, YM and JS are funded by (1) Research Council KU Leuven: C14/18/092 SymBioSys3; CELSA-HIDUCTION, (2) Innovative Medicines Initiative: MELLODY, (3) Flemish Government (ELIXIR Belgium, IWT: PhD grants, FWO 06260) and (4) Impulsfonds AI: VR 2019 2203 DOC.0318/1QUATER Kenniscentrum Data en Maatschappij. Computational resources and services used in this work were provided by the VSC (Flemish Supercomputer Center), funded by the Research Foundation - Flanders (FWO) and the Flemish Government – department EWI. We also gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research.

References

[1] Garrett B. Goh, Charles Siegel, Abhinav Vishnu, Nathan O. Hodas, and Nathan Baker. Chemception: A deep neural network with minimal chemistry knowledge matches the performance of expert-developed qsar/qspr models. *CoRR*, abs/1706.06689, 06 2017. URL `https://arxiv.org/abs/1706.06689`.

[2] Hongming Chen, Ola Engkvist, Yinhai Wang, Marcus Olivecrona, and Thomas Blaschke. The rise of deep learning in drug discovery. *Drug discovery today*, 23, 01 2018. URL `https://doi.org/10.1016/j.drudis.2018.01.039`.

[3] Francois Chollet. Xception: Deep learning with depthwise separable convolutions. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. URL `https://doi.org/10.1109/CVPR.2017.195`.

[4] Özgün Çiçek, Ahmed Abdulkadir, Soeren S. Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: Learning dense volumetric segmentation from sparse annotation. *CoRR*, abs/1606.06650, 2016. URL `http://arxiv.org/abs/1606.06650`.

[5] Connor W. Coley, Regina Barzilay, William H. Green, Tommi S. Jaakkola, and Klavs F. Jensen. Convolutional embedding of attributed molecular graphs for physical property prediction. *Journal of Chemical Information and Modeling*, 57(8):1757–1772, 2017. URL `https://doi.org/10.1021/acs.jcim.6b00601`.

[6] Arthur Dalby, James G. Nourse, W. Douglas Hounshell, Ann K. I. Gushurst, David L. Grier, Burton A. Leland, and John Laufer. Description of several chemical structure file formats used by computer programs developed at molecular design limited. *Journal of Chemical Information and Computer Sciences*, 32(3):244–255, 1992. URL `https://pubs.acs.org/doi/abs/10.1021/ci00007a012`.

[7] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alan Aspuru-Guzik, and Ryan P Adams. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in Neural Information Processing Systems 28*, pages 2224–2232. Curran Associates, Inc., 2015. URL `http://papers.nips.cc/paper/5954-convolutional-networks-on-graphs-for-learning-molecular-fingerprints.pdf`.

[8] Igor V. Filippov and Marc C. Nicklaus. Optical structure recognition software to recover chemical information: Osra, an open source solution. *Journal of Chemical Information and Modeling*, 49(3):740–743, 2009. URL `https://doi.org/10.1021/ci800067r`.

[9] Anna Gaulton, Anne Hersey, Michał Nowotka, A. Patrícia Bento, Jon Chambers, David Mendez, Prudence Mutowo, Francis Atkinson, Louisa J. Bellis, Elena Cibrián-Uhalte, Mark Davies, Nathan Dedman, Anneli Karlsson, María Paula Magariños, John P. Overington, George Papadatos, Ines Smit, and Andrew R. Leach. The ChEMBL database in 2017. *Nucleic Acids Research*, 45(Database issue):D945–D954, January 2017. ISSN 0305-1048. doi: 10.1093/nar/gkw1074. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5210557/`.

[10] Rafael Gómez-Bombarelli, David K. Duvenaud, José Miguel Hernández-Lobato, Jorge Aguilera-Iparraguirre, Timothy D. Hirzel, Ryan P. Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *CoRR*, abs/1610.02415, 2016. URL http://arxiv.org/abs/1610.02415.

[11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. URL https://doi.org/10.1109/CVPR.2016.90.

[12] Andreas Jansson, Eric J. Humphrey, Nicola Montecchio, Rachel M. Bittner, Aparna Kumar, and Tillman Weyde. Singing voice separation with deep u-net convolutional networks. In *ISMIR*, 2017. URL http://openaccess.city.ac.uk/id/eprint/19289.

[13] Steven Kearnes, Kevin McCloskey, Marc Berndl, Vijay Pande, and Patrick Riley. Molecular graph convolutions: moving beyond fingerprints. *Journal of computer-aided molecular design*, 30(8):595–608, 2016. URL https://doi.org/10.1007/s10822-016-9938-8.

[14] Greg Landrum. Rdkit: Open-source cheminformatics. URL http://www.rdkit.org. accessed on 03.06.2019.

[15] Yann LeCun, Yoshua Bengio, and Geoffrey E. Hinton. Deep learning. *Nature*, 521(7553): 436–444, 2015. URL https://doi.org/10.1038/nature14539.

[16] Noureddin M. Sadawi, Alan Sexton, and Volker Sorge. Chemical structure recognition: A rule based approach. *Proceedings of SPIE - The International Society for Optical Engineering*, 8297:32–, 01 2012. URL https://doi.org/10.1117/12.912185.

[17] Joe R. McDaniel and Jason R. Balmuth. Kekule: Ocr-optical chemical (structure) recognition. *Journal of Chemical Information and Computer Sciences*, 32(4):373–378, 1992. URL https://doi.org/10.1021/ci00008a018.

[18] Jungkap Park, Gus Rosania, Kerby A Shedden, Mandee Nguyen, Naesung Lyu, and Kazuhiro Saitou. Automated extraction of chemical structure information from digital raster images. *Chemistry Central journal*, 3:4, 03 2009. URL https://doi.org/10.1186/1752-153X-3-4.

[19] Douglas Pires, Tom Blundell, and David B Ascher. pkcsm: Predicting small-molecule pharmacokinetic and toxicity properties using graph-based signatures. *Journal of medicinal chemistry*, 58, 04 2015. URL https://doi.org/10.1021/acs.jmedchem.5b00104.

[20] O. Ronneberger, P.Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 9351 of *LNCS*, pages 234–241. Springer, 2015. URL http://lmb.informatik.uni-freiburg.de/Publications/2015/RFB15a.

[21] Evan Shelhamer, Jonathan Long, and Trevor Darrell. Fully convolutional networks for semantic segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(4):640–651, April 2017. URL https://doi.org/10.1109/TPAMI.2016.2572683.

[22] Jaak Simm, Adam Arany, Edward De Brouwer, and Yves Moreau. Graph Informer Networks for Molecules. *arXiv:1907.11318 [cs, stat]*, July 2019. URL http://arxiv.org/abs/1907.11318. arXiv: 1907.11318.

[23] Tim Soderberg. Chirality and stereoisomers. URL https://chem.libretexts.org/Bookshelves/Organic_Chemistry/Book%3A_Organic_Chemistry_with_a_Biological_Emphasis_(Soderberg)/Chapter_03%3A_Conformations_and_Stereochemistry/03.3%3A_Stereoisomerism_%E2%80%93_chirality%2C_stereocenters%2C_enantiomers. accessed on 03.06.2019.

[24] Joshua Staker, Kyle Marshall, Robert Abel, and Carolyn M. McQuaw. Molecular Structure Extraction from Documents Using Deep Learning. *Journal of Chemical Information and Modeling*, 59(3):1017–1029, March 2019. ISSN 1549-9596. doi: 10.1021/acs.jcim.8b00669. URL https://doi.org/10.1021/acs.jcim.8b00669.

[25] Aniko T. Valko and A. Peter Johnson. Clide pro: The latest generation of clide, a tool for optical chemical structure recognition. *Journal of Chemical Information and Modeling*, 49(4):780–787, 2009. URL https://doi.org/10.1021/ci800449t.

[26] Andreas Veit, Michael J Wilber, and Serge Belongie. Residual Networks Behave Like Ensembles of Relatively Shallow Networks. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 550–558. Curran Associates, Inc., 2016. URL http://papers.nips.cc/paper/6556-residual-networks-behave-like-ensembles-of-relatively-shallow-networks.pdf.

[27] David Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of Chemical Information and Computer Sciences*, 28(1):31–36, 1988. URL https://pubs.acs.org/doi/abs/10.1021/ci00057a005.

[28] Wikipedia: F1 score. F1 score, January 2020. URL https://en.wikipedia.org/w/index.php?title=F1_score&oldid=935689499. Page Version ID: 935689499.

[29] Zheng Xu, Sheng Wang, Feiyun Zhu, and Junzhou Huang. Seq2seq fingerprint: An unsupervised deep molecular embedding for drug discovery. In *Proceedings of the 8th ACM International Conference on Bioinformatics, Computational Biology,and Health Informatics*, ACM-BCB '17, pages 285–294, New York, NY, USA, 2017. ACM. URL http://doi.acm.org/10.1145/3107411.3107424.

[30] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *CoRR*, abs/1511.07122, 2016. URL https://arxiv.org/abs/1511.0712.

[31] Zhengxin Zhang, Qingjie Liu, and Yunhong Wang. Road extraction by deep residual u-net. *CoRR*, abs/1711.10684, 2017. URL http://arxiv.org/abs/1711.10684.