

Predicting Disease-Gene Associations using Cross-Document Graph-based Features

Hendrik ter Horst¹, Matthias Hartung¹, Roman Klinger^{1,2},
Matthias Zwick³, and Philipp Cimiano¹

¹ Semantic Computing Group, CITEC, Bielefeld University
{[hthors](mailto:hthors@techfak.uni-bielefeld.de),[mhartung](mailto:mhartung@techfak.uni-bielefeld.de),[cimiano](mailto:cimiano@techfak.uni-bielefeld.de)}@techfak.uni-bielefeld.de

² Institute for Natural Language Processing, University of Stuttgart
roman.klinger@ims.uni-stuttgart.de

³ Research Networking, Boehringer Ingelheim Pharma GmbH & Co. KG
matthias.zwick@boehringer-ingelheim.com

Abstract. In the context of personalized medicine, text mining methods pose an interesting option for identifying disease-gene associations, as they can be used to generate novel links between diseases and genes which may complement knowledge from structured databases. The most straightforward approach to extract such links from text is to rely on a simple assumption postulating an association between all genes and diseases that co-occur within the same document. However, this approach (i) tends to yield a number of spurious associations, (ii) does not capture different relevant types of associations, and (iii) is incapable of aggregating knowledge that is spread across documents. Thus, we propose an approach in which disease-gene co-occurrences and gene-gene interactions are represented in an RDF graph. A machine learning-based classifier is trained that incorporates features extracted from the graph to separate disease-gene pairs into valid disease-gene associations and spurious ones. On the manually curated *Genetic Testing Registry*, our approach yields a 30 points increase in F_1 score over a plain co-occurrence baseline.

Keywords: disease-gene associations, text mining, machine learning, biomedical literature, graph-based features

1 Introduction

Most current approaches in personalized medicine, irrespective of particular treatment modalities (e.g., small molecules, biologics, novel approaches like gene therapy), are centered around modulating a gene in order to modulate aspects of a disease [10]. Therefore, the detection of disease-gene links is an important starting point in drug discovery.

Text mining methods pose an interesting option for identifying disease-gene links, as they can be used to generate new target (and therefore treatment) hypotheses and, in combination with experimental data, support the prioritization of research aimed at the discovery of new drug targets. Until now, text mining methods for disease-gene associations mostly rely on the degree of textual

co-occurrence [19]. While those approaches are largely reliable in detecting well-known links for well-known diseases [1, 5, 18], such well-known links are of minor interest in drug research, as they do not support the discovery of new targets. Such novel targets are difficult to detect, as they often require the aggregation of evidence across individual documents; at the same time, they potentially shed light on yet unknown disease-gene links.

In this paper, we propose a classification model for predicting novel disease-gene associations from biomedical text. The model combines *private* (intra-document) as well as *public* (cross-document) knowledge as defined by Swanson et al. [21] in terms of features based on local co-occurrences within documents and relations between diseases and genes that have been aggregated across individual documents. In an evaluation against an existing database, we address the following research questions: (i) Can such a combined model outperform a purely co-occurrence-based approach? (ii) What is the impact of features measuring the connectivity of diseases and genes? (iii) What is the impact of graph-based features capturing interactions between genes across documents?

2 System Architecture




Fig. 1. Overview of system architecture

Our system architecture consists of the following components (cf. Fig. 1):

1. **Medline Corpus:** All analyses are done on Medline⁴, comprising a total of 21.5M abstracts.
2. **Information Extraction:** We rely on existing information extraction systems to identify disease names, genes/proteins and interactions between

⁴ <http://www.ncbi.nlm.nih.gov/pubmed>

them. For **disease recognition**, we use a state-of-the-art CRF tagger [11] that has been trained on the NCBI Disease Corpus [7]. Using the normalization procedure described in [22], disease names are reduced to a vocabulary of approx. 15K unique identifiers extracted from MeSH⁵ and OMIM⁶. For the **identification of genes/proteins and their interactions**, we rely on TEES [3], a state-of-the-art event extraction system that has been tailored to the detection of molecular interactions from biomedical text. All extracted genes are normalized using GeNo [24] and afterwards filtered by human genes using taxonomic information from EntrezGene [14].




Fig. 2. TEES representation of an example sentence from [13].

3. **Postprocessing:** The complexity of the interaction graphs produced by TEES (see Fig. 2) is incrementally reduced. Fig. 3 gives a running example of the individual steps described in the following:

- (a) **RDFication:** TEES graphs are broken down into binary relations by selecting all shortest paths that connect two proteins with respect to their semantic relation. These are represented as one RDF triple connecting two proteins. The path between the two proteins is serialized as a string and used as the name of an RDF property connecting both proteins.
- (b) **Simplification:** Semantic role information (i.e., CAUSE and THEME) is omitted. The direction of interactions is still captured in the directed edges of the graph.
- (c) **Generalization:** The TEES extraction scheme, originally consisting of 9 relations, was reduced to five relations after discussion with a domain expert: *Expression*, *Catabolism*, *Localization*, *Binding* and *Regulation*.
- (d) **Compression:** Consecutive occurrences of identical relations within a path signature are compressed by reducing them to one relation.
- (e) **Path joining:** To extract longer dependencies between genes/proteins as well, we join paths connecting two genes up to a distance of two edges. The join of the paths is serialized again, the above post-processing steps are applied and the results are stored as RDF triples. We refer to such serialized paths as *path signatures*.

⁵ <https://www.nlm.nih.gov/mesh/>

⁶ <http://omim.org>

(a)	HSP27 – THEME:Neg_reg:CAUSE:Pos_reg:THEME:Pos_reg:CAUSE – ActD HSP27 – THEME:Neg_reg:CAUSE:Pos_reg:THEME:Pos_reg:THEME – caspase3 ActD – CAUSE:Pos_reg:THEME – caspase3
(b)	HSP27 – Neg_reg:Pos_reg:Pos_reg – ActD HSP27 – Neg_reg:Pos_reg:Pos_reg – caspase3 ActD – Pos_reg – caspase3
(c)	HSP27 – Reg:Reg:Reg – ActD HSP27 – Reg:Reg:Reg – caspase3 ActD – Reg – caspase3
(d)	HSP27 – Reg3 – ActD HSP27 – Reg3 – caspase3 ActD – Reg – caspase3

Fig. 3. Example of the post-processing procedure applied to TEES events. Paths in (a) are equivalent to the original TEES output (cf. Fig. 2), (d) shows the result after post-processing, without path joining being applied.

- 4. Database:** The results of all the steps described above are stored in an RDF database, *Blazegraph*⁷.
- 5. Gene Classification System:** Given a disease as input, protein candidates are classified as to whether or not they interact with that disease. The classifier relies on features extracted for each pair of disease and protein. Being our main contribution, this component is described in the next section.

3 Gene Classification System

Our gene classification system takes a disease as input and predicts, for each gene in the database, whether it interacts with the given disease or not. The classifier is implemented as a Support Vector Machine relying on features that are extracted for each disease-gene pair. Seven feature groups are defined in total, which can be divided into *co-occurrence-based (CBF)* and *graph-based (GBF)* features, as described below.

Our notation is as follows: Let D be the set of all diseases and G the set of all genes in the database⁸ and P a vocabulary of predicates denoting semantic relations between them. Then, T denotes the set of all triples in the database, such that: $T \subset (D \times P \times G) = \{\langle d, p, g \rangle | p = coocc\} \cup \{\langle g, p, g' \rangle | p = interact\}$.

3.1 Co-occurrence-based Features

CBF features are based on the co-occurrence between diseases and genes. We consider $T_d \subset T = \{\langle d', coocc, g \rangle | d' = d\}$, the set of all genes co-occurring with a

⁷ <http://www.blazegraph.com>

⁸ D and G comprise all diseases and genes recognized during preprocessing the Medline corpus (cf. Section 2). This amounts to 7.640 diseases and 11.201 genes, in total.

particular disease d , and analogously $T_g \subset T = \{\langle d, coocc, g' \rangle | g' = g\}$. Moreover, $T_{dg} \subset T = \{\langle d', coocc, g' \rangle | d' = d, g' = g\}$ denotes the set of all co-occurrences of a particular disease d and a particular gene g .

Entropy. We compute the entropy $H(g)$ of a gene g in order to measure the specificity of g in terms of the diseases it co-occurs with. If g co-occurs with only a few specific diseases, this results in low entropy and high specificity. Co-occurrence with many diseases yields high entropy and low specificity. We compute the entropy of g as $H(g) = -\sum_d^D p(d|g) \cdot \log_2 p(d|g)$, where $p(d|g) = |T_{dg}|/|T_g|$. Analogously, we compute the entropy/specificity $H(d)$ of a disease d in terms of the genes it co-occurs with.

Co-occurrence Frequencies. This feature group combines relative co-occurrence frequencies of a disease-gene pair (d, g) :

$$Occ(d, g) = \frac{|T_{dg}|}{\max_{d' \in D} |T_{d'}|} \quad (1)$$

Besides the normalization given in Equation (1), two other alternatives are used. In all variants, $Occ(d, g)$ measures the strength of the connection of d and g . Ranging from 0 to 1, small values indicate a weak connection, whereas larger values indicate a strong connection.

Grades. This feature group consists of two features which capture a normalized frequency of triples that contain d or g , respectively: $Grade(d) = |T_d|/\max_{d' \in D} |T_{d'}|$ and $Grade(g) = |T_g|/\max_{g' \in G} |T_{g'}|$.

Odds Ratio is used to assess the degree of association between d and g :

$$Odds(d, g) = \frac{|T_{dg}| \cdot (|T| - |T_{dg}|)}{(|T_{dg}| - |T_d|) \cdot (|T_{dg}| - |T_g|)} \quad (2)$$

The higher $Odds(d, g)$, the stronger the association between d and g . $Odds(d, g) = 0$ can only be achieved if $|T_{dg}| = 0$.

TF-IDF. In order to assess the relevance of a gene g for a disease d , we apply the *tfidf* metric from information retrieval which takes term frequency (*tf*) and inverted document frequency (*idf*) into account [15]: $tfidf(d, g) = tf(d, g) \cdot idf(g)$. Considering a disease as a ‘‘bag of genes’’, $tf(d, g)$ is equivalent to $|T_{dg}|$, while $idf(g)$ can be computed in terms of Equation (3):

$$idf(g) = \log \left(\frac{|D|}{\sum_{d \in D} f(d, g)} \right), \text{ where } f(d, g) = \begin{cases} 1 & \text{if } |T_{dg}| > 0 \\ 0 & \text{else} \end{cases} \quad (3)$$

High values of $tfidf(d, g)$ indicate that g is mentioned frequently in the context of d , but still sufficiently specific to be informative for d , which we expect to be indicative of a relevant association between d and g .

3.2 Graph-based Features

In contrast to the previously described feature groups which take a disease and a gene into account, GBF features are calculated *independently* of a particular disease in that they are entirely based on the gene interaction graph, i.e., the set of triples $I \subset T = \{\langle g, interact, g' \rangle | g, g' \in G\}$.

Path Signatures. Each gene g is described in terms of a “bag of (outgoing) path signatures”, $S_{out}(g) \subset I = \{\langle g', interact, g'' \rangle | g = g'\}$, which have been constructed by joining individual edges in the gene interaction network (cf. Section 2, postprocessing step (3e)). We use *interact* as a placeholder for all predicates constructed in this process. The strength of interaction between a pair $\langle g, g' \rangle \in S_{out}(g)$ is weighted by the *tfidf* metric, expressing *tf*(g, g') and *idf*(g') analogously to the definitions in Section 3.1 (cf. Equation(3)). Path signatures encoding an important relation between g and g' in terms of a high *tfidf* value are considered useful for predicting novel disease-gene associations in cases where no direct evidence from co-occurrence relations is available yet.

Gene Connectivity. This feature group describes the connectivity of a gene within the graph. Analogously to $S_{out}(g)$ above, we define $S_{in}(g, l)$ and $S_{out}(g, l)$ as lists of all incoming and outgoing signatures of path length l , respectively. Further, L denotes the maximum path length. Based on these definitions, we count the number of incoming and outgoing signatures for each path length $1 \leq l \leq L$ separately, as given in (4), and by accumulation over all path lengths. In these features, higher values indicate a higher connectivity of g in the network.

$$Out_l(g) = \frac{|S_{out}(g, l)|}{\max_{g' \in G} |S_{out}(g', l)|} \quad In_l(g) = \frac{|S_{in}(g, l)|}{\max_{g' \in G} |S_{in}(g', l)|} \quad (4)$$

We also measure the ratio of outgoing and incoming signatures per gene in terms of $IORatio(g) = |S_{out}(g)|/|S_{in}(g)|$. If $IORatio(g) > 1$, g has a manipulating role in the network; otherwise, g tends to be manipulated by other genes.

4 Experimental Evaluation

4.1 Experimental Settings

Gold Standard. The *Genetic Testing Registry* (GTR; [20]) is a manually curated database for results from biomedical experiments, mostly at the intersection of Mendelian disorders and human genes. Our GTR dump contains 5,800 disease-gene associations built from 4,200 diseases and 2,800 genes.

Training and Testing Data are created from GTR as follows: All disease-gene associations in GTR are considered as positive examples. For each disease in GTR, we additionally generate the same amount of negative training examples by pairing the disease with genes that co-occur in Medline but are not attested in

GTR as a valid disease-gene association. The resulting data set is split into 80% used for training and 20% for testing. The training set contains 3,665 diseases with 1,781 negative and 1,884 positive examples (i.e., associated genes). The test set comprises 910 diseases with 440 positive and 470 negative examples.

Experimental Procedure. We train an SVM classifier using an RBF kernel [6] and apply grid search for meta-parameter optimization based on the LibSVM⁹ and WEKA¹⁰ toolkits. The trained model is applied to the task of predicting genes that are associated with a given disease. We evaluate the model on the GTR test set described above, reporting precision, recall and F₁ score.

4.2 Results and Discussion

Feature Group	Prec.	Rec.	F ₁	Feat. Combination	Prec.	Rec.	F ₁
Entropy	63.2	72.5	67.5	CBF	89.6	79.8	84.4
Co-occurrence	82.4	69.0	75.1	CBF+Connectivity	89.1	82.2	85.5
Grade	62.8	82.6	71.4	CBF+Conn+Best50Sig	89.4	84.9	87.1
Odds Ratio	77.9	43.8	56.1	CBF+Conn+Best100Sig	88.1	83.3	85.7
TF-IDF	85.3	62.8	72.3	CBF+Conn+Best200Sig	88.2	84.1	86.1
Path Signatures	66.6	75.6	70.8	Baseline	87.7	41.5	56.3
Connectivity	62.3	78.3	69.4				

Table 1. Evaluation results of classification models based on individual feature groups (left) and feature group combinations (right) on GTR testing data

Evaluation results are reported in Table 1. The left part displays classification performance of individual feature groups; in the right part, testing performance of feature combinations (as selected by cross-validation on the training data) are shown. The baseline refers to the performance of a single-feature classifier relying on co-occurrence counts as described in Equation (1).

The results clearly indicate a positive impact of both cooccurrence-based and graph-based features, as all feature combinations yield an increase over the baseline in both precision and recall. The CBF combination achieves highest overall precision, whereas connectivity and signature features improve recall (at the expense of slight losses in precision). As for path signatures, it is most effective to select only a small number of individual paths. We determine the best 50 path signature features based on information gain [12]. In the best-performing configuration (CBF+Conn+Best50Sig), our system outperforms the co-occurrence baseline by 1.7 points in precision and 43.4 points in recall.

Increasing the recall relative to the co-occurrence baseline is a key prerequisite towards our goal of discovering novel disease-gene associations. Given that

⁹ <http://www.csie.ntu.edu.tw/~cjlin/libsvm>

¹⁰ <http://www.cs.waikato.ac.nz/ml/weka/>

the GTR gold standard is relatively small and slightly biased towards Mendelian disorders, we subject the best-performing model to another evaluation in a practical use case, as described in the next section.

4.3 Case Study: Pulmonary Fibrosis

	Correct	Plausible Candidate	Incorrect
Count (Percentage)	77 (38.5%)	79 (39.5%)	44 (22.0%)

Table 2. Preliminary results of manual evaluation of 200 gene candidates predicted for pulmonary fibrosis.

In this experiment, our classification model was applied to the entire Medline corpus in order to predict genes related to *pulmonary fibrosis* (PF). The resulting hits were sorted by their corpus frequency and the 200 most frequent candidates were manually evaluated by a biomedical expert (who is not a PF researcher, though). Table 2 shows the preliminary results of this analysis: 38.5% of the predictions are unanimously correct, whereas only 22% are clear errors. The missing mass is due to candidates for which hints were found that the gene may be associated with PF through relevant mechanisms or pathways. Thus, these candidates constitute plausible hypotheses which need further investigation by a PF expert. Main sources of erroneous predictions are false co-occurrences (e.g., due to negation contexts) or false positives as produced by the gene recognition component. Some errors of the latter type may be eliminated by incorporating the filtering approach proposed by [9]. In sum, this analysis clearly shows that our system is capable of generating promising candidates worth further investigation.

5 Related Work

Three types of approaches have been proposed to tackle the problem of extracting explicitly mentioned disease-associated genes (DAGs) but also generating novel hypotheses from scientific publications. First, several authors extract DAGs from existing biomedical databases such as GeneSeeker [8] or PolySearch [4]. Piñero et al. developed DisGeNET [17], a database quantifying the degree of disease-gene associations by a combination of different sources of evidence, with textual co-occurrence being one of the main sources. Obviously, these approaches lack the ability to discover new target hypotheses. Second, text mining techniques have been considered as an alternative and are mostly based on textual co-occurrence (sentence or document-based). Such systems can be optimized on precision [5] or recall [18]. Al-Mubaid presents a technique using various information-theoretic concepts to support the co-occurrence-based extraction [1]. Third, a promising alternative to overcome mere textual co-occurrence is to aggregate knowledge

across single publications (cf. [23]) into larger interaction graphs, as we also do in our approach. Nevertheless, the knowledge extraction to build those interaction graphs often relies on text mining techniques and natural language processing methods as in the BITOLA system [2] or in the approaches of Wren et al. [26] and Wilkinson et al. [25]. Closely related to our approach is the one by Özgür et al. [16] who extract interaction paths from dependency networks and rely on graph centrality measures to rank proteins for a given disease. Contrary to our model, they do not use complex features extracted from the graph and do not combine different types of features.

6 Conclusions and Outlook

In this paper, we have presented a system and a model for predicting disease-gene associations from biomedical text, using a combination of features based on disease-gene co-occurrences and gene interactions that are represented in a graph database. In a classification experiment against a manually curated database used as gold standard, we were able to demonstrate the effectiveness of both types of features, outperforming a plain co-occurrence baseline by more than 30 points in F_1 score. Moreover, preliminary investigation of a practical use case from pharmaceutical industry suggests that almost 80% of the candidates predicted by our model are plausible and may support pharmaceutical researchers in hypothesis generation. In future work, we will carry out a more detailed evaluation of the case study and supplement our classification approach by a ranking model that not only separates positive and negative candidates but also reflects relative differences in these candidates' plausibility.

References

1. Al-Mubaid, H., Singh, R.K.: A new text mining approach for finding protein-to-disease associations. *American Journal of Biochemistry and Biotechnology* 1(3), 145 (2005)
2. Baud, R., et al.: Improving literature based discovery support by genetic knowledge integration. *The New Navigators: From Professionals to Patients* 95, 68 (2003)
3. Björne, J., Heimonen, J., Ginter, F., Airola, A., Pahikkala, T., Salakoski, T.: Extracting complex biological events with rich graph-based feature sets. In: *Proc. of the BioNLP Shared Task*. pp. 10–18 (2009)
4. Cheng, D., Knox, C., Young, N., Stothard, P., Damaraju, S., Wishart, D.S.: Polysearch: a web-based text mining system for extracting relationships between human diseases, genes, mutations, drugs and metabolites. *Nucleic acids research* 36(suppl 2), W399–W405 (2008)
5. Chun, H.W., Tsuruoka, Y., Kim, J.D., Shiba, R., Nagata, N., Hishiki, T., Tsujii, J.: Extraction of gene-disease relations from medline using domain dictionaries and machine learning. In: *Pacific Symposium on Biocomputing*. vol. 11, pp. 4–15 (2006)
6. Cortes, C., Vapnik, V.: Support-vector networks. *Machine learning* 20(3), 273–297 (1995)
7. Dogan, R.I., Lu, Z.: An improved corpus of disease mentions in pubmed citations. In: *Proc. of the BioNLP Workshop*. pp. 91–99 (2012)

8. van Driel, M.A., Cuelenaere, K., Kemmeren, P.P., Leunissen, J.A., Brunner, H.G.: A new web-based data mining tool for the identification of candidate genes for human genetic disorders. *European Journal of Human Genetics* 11(1), 57–63 (2003)
9. Hartung, M., Klinger, R., Zwick, M., Cimiano, P.: Towards gene recognition from rare and ambiguous abbreviations using a filtering approach. In: *Proc. of the BioNLP Workshop*. pp. 118–127 (2014)
10. Kisor, D.F., Kane, M.D., Talbot, J.N., Sprague, J.E.: *Pharmacogenetics, Kinetics and Dynamics for Personalized Medicine*. Jones & Bartlett (2014)
11. Klinger, R., Friedrich, C.M., Fluck, J., Hofmann-Apitius, M.: Named entity recognition with combinations of conditional random fields. In: *Proceedings of the second biocreative challenge evaluation workshop* (2007)
12. Kullback, S., Leibler, R.A.: On information and sufficiency. *The annals of mathematical statistics* pp. 79–86 (1951)
13. Ma, W., Teng, Y., Hua, H., Hou, J., Luo, T., Jiang, Y.: Upregulation of heat shock protein 27 confers resistance to actinomycin d-induced apoptosis in cancer cells. *FEBS* 280, 4612–4624 (2013)
14. Maglott, D., Ostell, J., Pruitt, K.D., Tatusova, T.: Entrez gene: gene-centered information at ncbi. *Nucleic acids research* 39(suppl 1), D52–D57 (2011)
15. Manning, C.D., Raghavan, P., Schütze, H.: *Introduction to Information Retrieval*. Cambridge University Press (2008)
16. Özgür, A., Vu, T., Erkan, G., Radev, D.R.: Identifying gene-disease associations using centrality on a literature mined gene-interaction network. *Bioinformatics* 24(13), i277–i285 (2008)
17. Piñero, J., Queralt-Rosinach, N., Bravo, À., Deu-Pons, J., Bauer-Mehren, A., Baron, M., Sanz, F., Furlong, L.I.: Disgenet: a discovery platform for the dynamical exploration of human diseases and their genes. *Database* 2015, bav028 (2015)
18. Pletscher-Frankild, S., Pallejà, A., Tsafou, K., Binder, J.X., Jensen, L.J.: Diseases: Text mining and data integration of disease–gene associations. *Methods* (2014)
19. Quan, C., Ren, F.: Gene–disease association extraction by text mining and network analysis. In: *Proceedings of the 5th International Workshop on Health Text Mining and Information Analysis (Louhi)@ EACL*. pp. 54–63 (2014)
20. Rubinstein, W.S., Maglott, D.R., Lee, J.M., Kattman, B.L., Malheiro, A.J., Ovetzky, M., Hem, V., Gorelenkov, V., Song, G., Wallin, C., et al.: The nih genetic testing registry: a new, centralized database of genetic tests to enable access to comprehensive information and improve transparency. *Nucleic acids research* p. gks1173 (2012)
21. Swanson, D.R., Smalheiser, N.R.: Undiscovered public knowledge: A ten-year update. In: *KDD*. pp. 295–298 (1996)
22. ter Horst, H.: Ranking of disease gene associations from large corpora of scientific publications. M.Sc. Thesis, Bielefeld University (5 2015)
23. Van Landeghem, S., Björne, J., Wei, C.H., Hakala, K., Pyysalo, S., Ananiadou, S., Kao, H.Y., Lu, Z., Salakoski, T., Van de Peer, Y., et al.: Large-scale event extraction from literature with multi-level gene normalization. *PLoS One* 8(4), e55814 (2013)
24. Wermter, J., Tomanek, K., Hahn, U.: High-performance gene name normalization with geno. *Bioinformatics* 25(6), 815–821 (2009)
25. Wilkinson, D.M., Huberman, B.A.: A method for finding communities of related genes. *Proc. of the National Academy of Science* 101(suppl 1), 5241–5248 (2004)
26. Wren, J.D., Bekeredjian, R., Stewart, J.A., Shohet, R.V., Garner, H.R.: Knowledge discovery by automated identification and ranking of implicit relationships. *Bioinformatics* 20(3), 389–398 (2004)